In the Quest for Scale-optimal Mappings

Vladimir Garanzha, Igor Kaporin, Liudmila Kudryavtseva, Francois Protais, Dmitry Sokolov

In the Quest for Scale-Optimal Mappings

VLADIMIR GARANZHA, IGOR KAPORIN, and LIUDMILA KUDRYAVTSEVA, Dorodnicyn Computing Center FRC CSC RAS, Moscow Institute of Physics and Technology, Moscow, Russia
FRANCOIS PROTAIS and DMITRY SOKOLOV, Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France

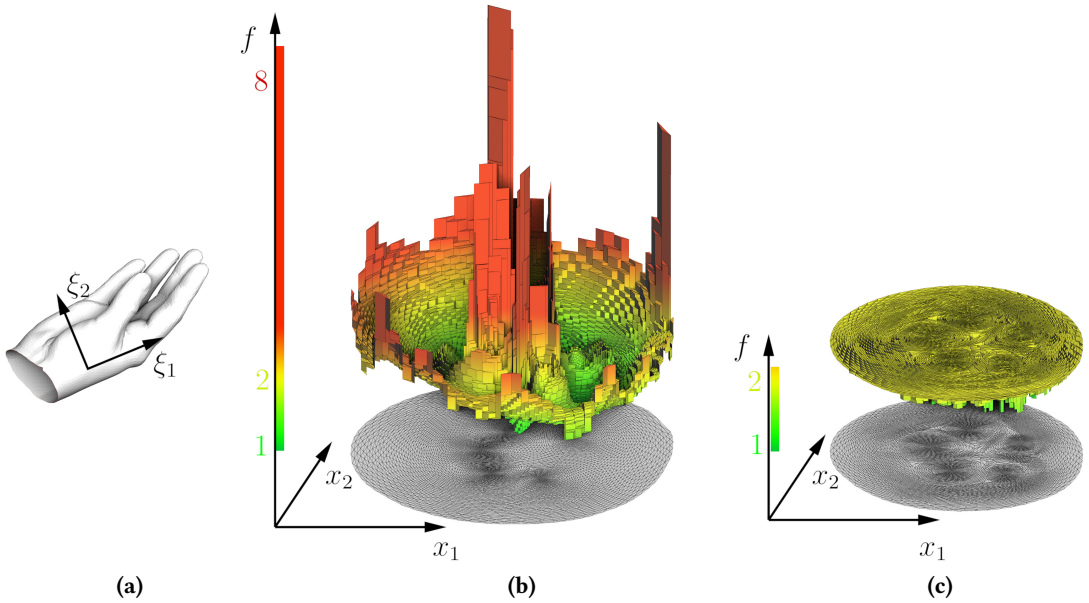

Fig. 1. Distortion distribution plotted over two flattenings of the input triangulated surface (a). Here, we flatten a 2-manifold surface parameterized by a 2D vector $\vec{\xi}$, where a flattening is represented by a 2D vector function $\vec{x}(\vec{\xi})$. We have defined a piecewise constant scalar distortion measure $f(x_1, x_2)$, with plots provided in (b) and (c). To enhance visual perception, we have depicted the piecewise constant function as a collection of colored prisms, with linear interpolation of colors for vertical facets. Elastic flattening (b) sacrifices the quality of several triangles to minimize average distortion. On the other hand, quasi-isometric stiffening (c) evenly distributes distortion over the domain, effectively minimizing the peaks of distortion.

Optimal mapping is one of the longest-standing problems in computational mathematics. It is natural to measure the relative curve length error under map to assess its quality. The maximum of such error is called the quasi-isometry constant, and its minimization is a nontrivial max-norm optimization problem. We present a physics-based quasi-isometric stiffening (QIS) algorithm for the max-norm minimization of hyperelastic distortion.

QIS perfectly equidistributes distortion over the entire domain for the ground truth test (unit hemisphere flattening) and, when it is not possible, tends to create zones where all cells have the same distortion. Such zones correspond to fragments of elastic material that became rigid under stiffening, reaching the deformation limit. As such, maps built by QIS are related to the de Boor equidistribution principle, which asks for an integral of a certain error indicator function to be the same over each mesh cell.

Under certain assumptions on the minimization toolbox, we prove that our method can build, in a finite number of steps, a deformation whose maximum distortion is arbitrarily close to the (unknown) minimum. We performed extensive testing: on more than 10,000 domains QIS was reliably better than the competing methods. In summary, we reliably build 2D and 3D mesh deformations with the smallest known distortion estimates for very stiff problems.

CCS Concepts: • **Computing methodologies** → **Mesh models**.

Additional Key Words and Phrases: Parameterization, injective mapping, mesh untangling, bounded distortion, quality mapping

Authors' addresses: Vladimir Garanzha; Igor Kaporin; Liudmila Kudryavtseva, Dorodnicyn Computing Center FRC CSC RAS, and Moscow Institute of Physics and Technology, Moscow, Russia; Francois Protais; Dmitry Sokolov, Université de Lorraine, and CNRS, and Inria, and LORIA, F-54000 Nancy, France.

1 INTRODUCTION

Origins of the problem. The construction of optimal deformations is one of the central themes in mesh generation research. While the principle of optimality itself is still controversial and problem-dependent, we follow the general formulation suggested by John Milnor in his paper "A Problem in Cartography" [Milnor 1969]. He proposed to analyze the relative change of the length of a curve in a map, referred to as "the scale" in terms of cartography. Optimal maps minimize the ratio of the maximum to minimum scale over the entire map. The idea itself is not new and dates back to P. L. Chebyshev and his famous talk for the Imperial Russian Geographic Society [Chebyshev 1856]. Chebyshev formulated the theorem stating that a conformal parameterization of a disc-like domain on the sphere

is scale-optimal if the unit scale is prescribed on its boundary. He constructed a conformal map for the European part of the Russian Empire, reducing the scale error to 1/50, providing a significant advantage over all other maps in use at the time (see Fig. 2 for an illustration). The first proof of Chebyshev's theorem was provided by Gravé in 1896 [Gravé 1911], and a modern rigorous proof was given by Milnor, who also proved that equidistant azimuthal flattening on the hemisphere is scale-optimal. Milnor's solution for the hemisphere and the cone parametrization by Bonk and Lang [Bonk and Lang 2003] are examples of the very few scale-optimal parameterizations known in closed form.

The goal. In this paper, our goal is to construct a deformation of a two- or three-dimensional domain Ω, possibly subject to certain constraints. We consider a map $\vec{x}(\vec{\xi}) : \Omega \to \Omega_x$, where $\Omega, \Omega_x \subset \mathbb{R}^d$, and d represents the dimension (2 or 3), with the arrow denoting a d-dimensional vector. Let us denote the Jacobian matrix of the map as J. By having the metric tensors $G_\xi(\vec{\xi})$ and $G_x(\vec{x})$, we can measure the lengths of curves in the domains Ω and Ω_x.

The idea is to study how the length of a curve in Ω relates to the length of its image in Ω_x under the action of the map $\vec{x}(\vec{\xi})$. For a simple curve defined by a 1D parameterization $\vec{\xi}(q) \in \Omega, 0 \le q \le Q$, its length can be found as follows:

$$L_\xi = \int_0^Q (\dot{\vec{\xi}}^\top G_\xi \dot{\vec{\xi}})^{\frac{1}{2}}\, dq.$$

The length of the curve's image $\vec{x}(\vec{\xi}(q))$ is defined as (recall that $\dot{\vec{x}} = J\dot{\vec{\xi}}$):

$$L_x = \int_0^Q (\dot{\vec{\xi}}^\top J^\top G_x J \dot{\vec{\xi}})^{\frac{1}{2}}\, dq.$$

A map $\vec{x}(\vec{\xi})$ is called scale-bounded (or quasi-isometric) if the following inequality holds for any curve $\vec{\xi}(q)$:

$$\frac{1}{K} L_\xi < L_x < K L_\xi,$$

where K is called the quasi-isometry constant. Note that this inequality makes sense for general rectifiable curves. For a map regular enough (say, piecewise differentiable one), we can reformulate this as a local matrix inequality:

$$\frac{1}{K^2} G_\xi < J^\top G_x J < K^2 G_\xi. \tag{1}$$

Our goal is to build a map with the lowest possible quasi-isometry constant K, known as a scale-optimal map. When this constant is equal to 1, the map preserves distances, reducing it to rigid body transformations.

[1]We were unable to find the original drawings by Chebyshev himself. For this illustration, we took a Chebyshev projection built by Dinchenko [Dinchenko 1938]. He approximated the boundary of the Soviet Union in the Mercator projection by an ellipse and computed an approximate solution for the Chebyshev map using a 6th-order polynomial expansion in complex variables. Note that the paper contains a few typographical errors; in our experiments, we had to flip the sign of the imaginary part of the 6th-order harmonic.

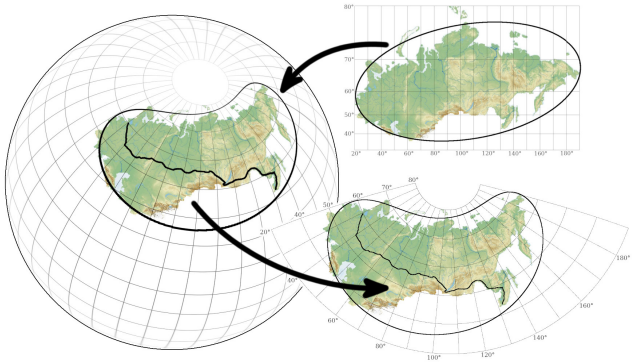

Fig. 2. To build the best conformal map (bottom right), in 1856 Chebyshev proposed[1] to prescribe the zero distortion (isometry) constraint along a closed curve. Here, the curve is chosen to be an ellipse in the Mercator projection (top right, 1:2 scale). If we measure the length of the Saint Petersburg–Vladivostok railway in the Chebyshev map (after normalization), we obtain less than 1.2% of error, i.e., 103 km over almost 9000 kilometers of the railroad.

After normalization of the map from Fig. 2, the closed curve delimits a domain where the scale varies from ≈ 1.027 (at the boundary) to ≈ 0.9734, resulting in the constant $K \approx 1.027$.

Approach. In this paper we solve a proxy problem approximating scale-optimal maps, refer to App. B.1 for a brief study of the relationship between the proxy and the original problems. We adopt a physical analogy with elastic deformations, wherein we construct a hyperelastic material that has an infinite stored energy if the pointwise distortion of hyperelastic deformation exceeds a certain threshold. As such, all admissible elastic deformations with finite energy are scale-bounded in nature.

Our contributions. In this paper, we formulate a continuation problem for the deformation threshold. By solving a series of fixed-threshold optimization problems, we obtain a physics-based quasi-isometric stiffening (QIS) algorithm for the max-norm minimization of hyperelastic distortion which is guaranteed to minimize certain upper estimate for quasi-isometry constant K. Under the assumption that there exists a minimization toolbox that is able to minimize a function by a whatever small, but constant rate, we prove that QIS builds a deformation arbitrarily close to the (unknown) optimum bound in a finite number of steps. Note that while in theory this assumption can only be satisfied by a global optimization solver, in practice we have not observed cases of false local minima when constructing lowest distortion mappings, even if counterexamples probably exist.

Our **parameter-free** approach is a discretization of a well-posed variational scheme, and it has an advantage that type, size and quality of mesh elements in the deformed object have a weak influence on the computed deformation. Our method is numerically stable, the positive definite part of the Hessian matrix is spectrally equivalent to the finite element Laplacian.

2 STATE OF THE ART

While various matrix-based optimization tools could be used to solve for scale-optimal mapping, most state-of-the-art methods use finite hyperelasticity functionals as a proxy.

2.1 Elastic deformations

Mathematical elasticity essentially involves minimizing the deviation from the isometric deformation state on average [Ciarlet 1988]. The idea is to consider a mesh as an elastic material, where the stored energy of deformation can be measured as an integral of a distortion measure. Therefore, the goal is to minimize the stored energy of deformation[2].

To build a good map $\vec{x}(\vec{\xi})$, we can solve the following variational problem:

$$\underset{\vec{x}(\vec{\xi})}{\arg\min} \int_{\Omega} f(J)\, d\xi, \tag{2}$$

where f is a measure of distortion, and

$$d\xi = \sqrt{\det G_{\xi}}\, d\xi_1 \cdot \cdots \cdot d\xi_d$$

is the volume differential that takes into account the metric in Ω.

John Ball in [1976] formulated the "polyconvexity" condition, which guarantees that the Euler-Lagrange equations for optimal deformation are elliptic. In formal terms, a function $f(J)$ is said to be polyconvex if it can be expressed as a convex function of the minors of the matrix J. In practice, ellipticity implies smoothness of deformations and continuous dependence on variation of constraints. It is not surprising that this choice has resulted in efficient engineering mesh generation algorithms with sound theoretical foundations [Jacquotte 1988; Rumpf 1996].

To the best of our knowledge, the first attempt to formulate a variational principle minimizing the deviation from isometry is due to G. B. Airy. In [1861], he suggested minimizing the following functional:

$$\underset{\vec{x}(\vec{\xi})}{\arg\min} \int_{\Omega} f_{Airy}(J)\, d\xi, \qquad f_{Airy}(J) := \left(\frac{\sigma_1}{\sigma_2} - 1\right)^2 + (\sigma_1\sigma_2 - 1)^2, \tag{3}$$

where σ_1 and σ_2 are the singular values of the Jacobian matrix J. Airy's "Balance-of-Errors" functional is the precursor of variational methods that trade shape against area distortion and vice versa. Interestingly, Airy refers to each term in the distortion measure f_A as "evil", while their sum is called "total evil". Assuming that $\sigma_i \approx 1$ and neglecting higher-order terms with respect to $\epsilon_i = 1 - \sigma_i$, Airy derived a simplified quadratic energy:

$$f_{Airy}^{ARAP}(J) := (\sigma_1 - 1)^2 + (\sigma_2 - 1)^2,$$

which corresponds exactly to the As-Rigid-as-Possible [Sorkine and Alexa 2007] energy.

[2]Refer to App. B.1 for a discussion on the fact that minimizing a distortion measure implies minimizing the upper bound for the quasi-isometry constant.

2.2 Constrained optimization

The equidistribution principle says that ideally the distortion is to be constant over the domain. This tends to happen if we minimize the maximum distortion (see Fig. 1). However, the generic "Balance-of-Errors" formulation (2) only allows us to minimize distortion on average. The problem of constructing deformations with bounded distortion has a long-standing history in elasticity research and goes back to Prager's work on "Ideal locking materials" [1957]. Nowadays, this problem is referred to as the "stiffening" of elastic material and is generally formulated as a set of nonlinear constraints on the Jacobian matrix of elastic deformation [Ciarlet and Necas 1985].

There is a large family of methods that try to take into account deformation bounds directly. For example, Sorkine et al. [2002] propose to lay triangles in a plane in a greedy manner without exceeding a user-specified distortion bound. Obviously, the mesh is cut during the procedure, and since it is possible to lay individual triangles without any distortion, the method succeeds.

Fu et al. [2016] propose to enforce the distortion constraints with a penalty method, leading to conflicts between multiple terms in the energy to minimize. Kovalsky et al. [2014] also attempt to control singular values of the Jacobian. This problem has a highly nonlinear and non-convex nature, and to do so, they resort to semidefinite programming under linear matrix inequality constraints. In a later work, Kovalsky et al. [2015] and Su et al. [2019] alternate between energy optimization and a non-trivial projection to the highly non-convex set of constraints. Lipman [2012] as well as Chien et al. [2016] and Levi and Zorin [2014] formulate the problem as a second-order cone programming, relying on elaborate commercial solvers such as MOSEK [Andersen and Andersen 2000]. In a similar manner, Aigerman and Lipman [2013] propose to solve a series of quadratic programs with linear constraints to build a bounded distortion map.

All these methods use generic optimization toolboxes, and while the approach may work reasonably well in practice, it is hard to obtain any guarantees.

2.3 Unconstrained optimization

Fortunately, it is possible to formulate the stiffening problem as an unconstrained minimization by introducing a weighting function that penalizes large values of distortion. Having a distortion measure $f(J)$, we can try to minimize the following energy:

$$\underset{\vec{x}(\vec{\xi})}{\arg\min} \int_{\Omega} w f(J)\, d\xi, \tag{4}$$

where w is a weight function. Typically, w is chosen to be large in the regions where small values of $f(J)$ are required. With a clever choice of w, this general weighted formulation can be used to control the pointwise behavior of the spatial distribution of the distortion measure.

Bommes et al. [2009] used such a weighting function in a heuristic procedure for mesh untangling. If an adaptation metric is prescribed in the computational domain, one can compute the weight function w according to this metric [Ivanenko 2000]. It is possible to penalize large values of local distortion by a power law enhancement. Fang

et al. [2021] and Garanzha et al. [2019] use the following definition: $w := f^p(J), p > 0$. Fu et al. [2015] propose a solution very similar in spirit to (4), using an exponential function to penalize the maximum distortion and minimizing $\int \exp(s \cdot f(J))$. They propose to use $s = 5$ as a rule of thumb.

While all these methods penalize large values of distortion, they do not offer a way to compute a deformation with a prescribed bound of distortion. Let us consider another definition of the weighting function proposed in [Garanzha 2000]:

$$w := \begin{cases} \dfrac{1}{1 - t f(J)}, & f(J) < \frac{1}{t} \\ +\infty, & f(J) \geq \frac{1}{t} \end{cases} \tag{5}$$

where $0 < t < 1$ is a distortion bound parameter. If f is chosen in a way that $f(J) \geq 1$, then a finite solution for Prob. (4) implies that $f(J) < \frac{1}{t}$. The resulting functional has an infinite barrier (refer to §3.3 for more details) on the boundary of the set of quasi-isometric deformations, thus solving the problem of quasi-isometric map generation for general domains.

Intuitively, it corresponds to a deformation of an object made from a quasi-isometric hyperelastic material. The material offers control over the distortion bound built directly into the definition of the deformation energy density. In this way, when a local measure of deformation exceeds a certain threshold, the elastic material becomes infinitely stiff. An invertibility theorem for deformations of this material was established both in 2D and 3D cases [Garanzha et al. 2014].

To optimize for the distortion bound, the stiffening threshold t is introduced as a parameter, and the max-norm optimization problem for deformation is formulated as a continuation problem for polyconvex functional (minimization of the stiffening threshold, or alternatively, maximization of the quality threshold). Unfortunately, until now a robust strategy of stiffening parameter choice was lacking.

In our paper we follow the formulation (5). Having carefully designed a strategy of choice for the stiffening parameter t, we obtain lowest distortion maps with our quasi-isometric stiffening (QIS) algorithm (Alg. 2 + Eq. (20)).

By coupling our technique with [Garanzha et al. 2021], we obtain a complete unified mapping pipeline. We start from an arbitrary initial deformation, untangle the mesh, thus minimizing mean distortion, and finally we minimize the maximum distortion. Both parts of the pipeline build upon the same ideas, and require only a linear solver [Hestenes and Stiefel 1952] for positive definite matrices (if Newton minimization is adopted) or a L-BFGS solver [Zhu et al. 1997] for a quasi-Newton scheme. We show that despite being quite different problems, foldover-free and bounded distortion mapping problems can be solved with very similar approaches (compare Alg. 1 and Alg. 2).

3 VARIATIONAL FORMULATION FOR UNTANGLING AND DISTORTION MINIMIZATION

This section provides a necessary primer on elastic deformations. After revisiting all the necessary concepts, we present our contributions on lowest distortion guarantees (§4) and testing results (§5).

3.1 Choice of elastic material

While there are multiple possible choices for the distortion measure (refer to App. B.1), for our meshes, we chose a material whose stored energy of deformation $\int_\Omega f(J)\, d\xi$ can be measured with density f defined as follows [Garanzha 2000]:

$$f(J) := (1 - \vartheta) f_s(J) + \vartheta f_v(J), \tag{6}$$

where shape distortion is defined as [Hormann and Greiner 2000]:

$$f_s(J) := \begin{cases} \dfrac{1}{d} \dfrac{\operatorname{tr} J^\top J}{(\det J)^{\frac{2}{d}}}, & \det J > 0 \\ +\infty, & \det J \leq 0 \end{cases} \tag{7}$$

and volumetric distortion is defined by:

$$f_v(J) := \begin{cases} \frac{1}{2}\left(\det J + \dfrac{1}{\det J}\right), & \det J > 0 \\ +\infty, & \det J \leq 0 \end{cases} \tag{8}$$

The functions $f_s(J)$ and $f_v(J)$ have concurrent goals: one tries to preserve angles, and the other tries to preserve volumes. The parameter ϑ serves as a trade-off parameter.

Note that the function f is not convex, but polyconvex. The notion of polyconvexity is a generalization of the notion of convexity for functions defined on spaces of matrices. In general, a function $\varphi(J) : \mathbb{R}^{d \times d} \to \mathbb{R} \cup + \{\infty\}$ is said to be polyconvex [Ball 1976] if there exists a convex function $\Phi(\#J)$ such that $\varphi(J) = \Phi(\#J)$, where $\#J$ denotes the set of all minors of J. For instance, any polyconvex function φ is rank-one convex, i.e.,

$$\varphi((1 - \epsilon)J + \epsilon(J + \delta J)) \leq (1 - \epsilon)\varphi(J) + \epsilon\varphi(J + \delta J),$$

where $\operatorname{rank} \delta J = 1$. Therefore, it satisfies the Hadamard-Legendre conditions (ellipticity conditions for the Euler-Lagrange equation of variational Prob. (2)):

$$\sum_{i,k,j,m=1}^{d} \frac{\partial^2 \varphi}{\partial(J)_{ij}\, \partial(J)_{km}} p_i p_k q_j q_m \geq 0$$

for arbitrary vectors $\vec{p}, \vec{q} \in \mathbb{R}^d$.

For our choice of material, it suffices to consider a narrower family of polyconvex functions where the extended function Φ is a convex function of $d^2 + 1$ arguments and $f(J) = \Phi(J, \det J)$. In summary, the energy density (6) is a polyconvex function, and therefore, the Euler-Lagrange equations for stored energy are elliptic, making them well-suited for numerical optimization.

3.2 Untangling

With a slight abuse of notation, the energy density (6) can be rewritten as follows:

$$f := (1-\vartheta)\frac{1}{d}\frac{\operatorname{tr} J^\top J}{(\max(0,\det J))^{\frac{2}{d}}} + \vartheta\frac{1}{2}\frac{1+\det^2 J}{\max(0,\det J)}$$

Note that if an initial guess is not admissible (has inverted elements), then the function is not defined. To overcome this problem, we can introduce the function $\chi(D,\varepsilon)$ that serves as a regularization for $\max(0,D)$ in the denominator:

$$\chi(D,\varepsilon) := \frac{D+\sqrt{\varepsilon^2+D^2}}{2} \tag{9}$$

When ε tends to zero, $\chi(\varepsilon,D)$ tends to $\max(0,D)$. Then, the energy density can be regularized following the ideas from [Garanzha and Kaporin 1999]:

$$f_\varepsilon(J) := (1-\vartheta)\frac{1}{d}\frac{\operatorname{tr} J^\top J}{(\chi(\det J,\varepsilon))^{\frac{2}{d}}} + \vartheta\frac{1}{2}\left(\chi(\det J,\varepsilon) + \frac{1}{\chi(\det J,\varepsilon)}\right). \tag{10}$$

Finally, Prob. (2) can be rewritten as follows:

$$\lim_{\varepsilon\to 0^+} \operatorname*{arg\,min}_{\vec{x}(\vec{\xi})} \int_\Omega f_\varepsilon(J)\, d\xi \tag{11}$$

This formulation suggests an algorithm: build a decreasing sequence of the $\varepsilon^k \to 0$, and for each value ε^k solve an optimization problem. In other words, Prob. (11) offers a way of getting rid of foldovers by solving a continuation problem with respect to the parameter ε.

The simplest way to discretize Prob. (11) is with first-order FEM: the map $\vec{x}$ is piecewise affine with the Jacobian matrix J being piecewise constant, and can be represented by the coordinates of the vertices in the computational domain $\{\vec{x}_i\}_{i=1}^{\#V}$. Let us denote the vector of all variables as $X := \left(\vec{x}_1^\top \dots \vec{x}_{\#V}^\top\right)^\top$.

A discretization of Prob. (11) can be written as follows:

$$\lim_{\varepsilon\to 0^+} \operatorname*{arg\,min}_X F(X,\varepsilon), \tag{12}$$
$$\text{where}\quad F(X,\varepsilon) := \sum_{k=1}^{\#T} f_\varepsilon(J_k)\operatorname{vol}(T_k),$$

Here, $\#V$ is the number of vertices, $\#T$ is the number of simplices, J_k is the Jacobian matrix for the k-th simplex, and $\operatorname{vol}(T_k)$ is the signed volume of the simplex T_k in the parametric domain.

Function $F(X,0)$ has an impenetrable infinite barrier on the boundary of the

set of meshes with positive cell volumes[3]:

$$\frac{\operatorname{vol}(\vec{x}(T_k))}{\operatorname{vol}(T_k)} > 0, \quad k = 1,\dots,\#T \tag{13}$$

This is a finite-dimensional approximation of the set $\det J > 0$.

Prob. (12) can be solved with Alg. 1. Untangling in Prob. (12) is guaranteed because [Garanzha et al. 2021] build a decreasing sequence $\varepsilon^k \to 0$ that forces the mesh to fall into the feasible set (13) of untangled meshes. With some assumptions on the minimization toolbox chosen, untangling is guaranteed to succeed in a finite number of steps.

3.3 Stiffening idea

In addition to untangling, by solving Prob. (12), we minimize the **average** distortion of a map. Let us for a short while suppose that we have a user-specified distortion bound t^*. How can we build a map whose **maximum** distortion respects the bound?

Consider the following variational problem related to the construction of deformations with a prescribed quality $0 \le t^* < 1$ [Garanzha 2000]:

$$\operatorname*{arg\,min}_{\vec{x}(\vec{\xi})} \int_\Omega \frac{f(J)}{1-t^* f(J)}\, d\xi \tag{14}$$

Recall that $f(J) \ge 1$, so for this integral to be finite, a necessary condition is:

$$f(J) < \frac{1}{t^*} \tag{15}$$

A finite-dimensional approximation of Prob. (14) can be written as follows:

$$\lim_{t^*-t\to+0} \operatorname*{arg\,min}_X W(X,t), \tag{16}$$
$$\text{where}\quad W(X,t) := \sum_{k=1}^{\#T} \frac{f(J_k)}{1-t f(J_k)}\operatorname{vol}(T_k),$$

Recall that to optimize for the average distortion, we needed a special untangling procedure to reach the

set (13) of foldover-free deformations. Here, we need to find a deformation that belongs to the set of deformations with a prescribed quality t^*:

$$f(J_k) < \frac{1}{t^*}, \quad k = 1,\dots,\#T \tag{17}$$

To do so, it is possible to use an approach that shares the spirit with untangling. It is important to note that a solution of Prob. (12) corresponds to a solution of Prob. (16) with $t^* = 0$, i.e., when no bound on the maximum deformation is imposed. But then, having reached the set (13), we can build an increasing sequence of $t^k \to t^*$ to contract the set until the mesh falls into the set (17) of deformations with a prescribed quality t^*.

Alg. 2 summarizes the optimization procedure, note how closely it is related to Alg. 1. While the general idea was published more than 20 years ago, until now, it remained unclear how to build this sequence $\{t^k\}$, and this constitutes the main theoretical contribution of the present paper.

[3] As a side note, this set has a quite complicated structure. For the k-th simplex, $\operatorname{vol}(\vec{x}(T_k))$ is a polylinear function of the coordinates of its vertices, hence each term in (13) defines a non-convex set. One can hardly expect that the intersection of the sets in (13) would result in a convex domain. Moreover, Ciarlet [Ciarlet and Geymonat 1982] has proved that the barrier property and convexity of the density of deformation energy are incompatible. Fortunately, barrier distortion measures can be polyconvex, as shown by J. Ball [Ball 1976].

Input: X^0 *// arbitrary initial guess (vector of size $\#V \times d$)*
Output: X *// final foldover-free map (vector of size $\#V \times d$)*
1: $k \leftarrow 0$;
2: **repeat**
3: compute ε^k; *// decreasing sequence*
4: $X^{k+1} \leftarrow \arg\min_X F(X, \varepsilon^k)$;
5: $k \leftarrow k + 1$;
6: **until** $\min_{t \in 1 \ldots \#T} \det J_t^k > 0$ **and** $F(X^k, \varepsilon^k) > (1 - 10^{-3})\, F(X^{k-1}, \varepsilon^{k-1})$
7: $X \leftarrow X^k$;

Alg. 1. Computation of a foldover-free map

Input: X^0 *// untangled initial guess (vector of size $\#V \times d$)*
Output: X *// final bounded distortion map (vector of size $\#V \times d$)*
1: $k \leftarrow 0$;
2: **repeat**
3: compute t^k; *// increasing sequence, Eq. (20)*
4: $X^{k+1} \leftarrow \arg\min_X W(X, t^k)$;
5: $k \leftarrow k + 1$;
6: **until** $W(X^k, t^k) > (1 - 10^{-3})\, W(X^{k-1}, t^{k-1})$
7: $X \leftarrow X^k$;

Alg. 2. Quasi-isometric stiffening (QIS)

4 LOWEST DISTORTION MAPPING

While we assume that parameter t^* exists, fortunately we are not obliged to know it to make the QIS algorithm work. It is an important advantage over optimization algorithms which use a prescribed distortion bound like LBD [Kovalsky et al. 2015], since in practice even rough estimates of this bound are not available. Essentially, the QIS algorithm by itself can serve as a distortion bound estimation tool for problems of any complexity.

This section provides our main result, namely a way to build an increasing sequence $\{t^k\}$ that allows us to effectively contract the feasible set until we reach the goal. Untangling and stiffening are very closely related, so let us first restate the main result of [Garanzha et al. 2021, Theorem 1]. It will allow us to highlight the similarity between the approaches.

Theorem 1. *Let us suppose that the feasible set of untangled meshes* (13) *is not empty. We also suppose that for solving $X^{k+1} \leftarrow \arg\min_X F(X, \varepsilon^k)$ we have a minimization algorithm satisfying the following efficiency conditions for some $0 < \sigma < 1$:*
For each iteration k,

- ***either*** *the essential descent condition holds*

$$F(X^{k+1}, \varepsilon^k) \le (1 - \sigma) F(X^k, \varepsilon^k), \tag{18}$$

- ***or*** *the vector X^k satisfies the quasi-minimality condition:*

$$\min_X F(X, \varepsilon^k) > (1 - \sigma) F(X^k, \varepsilon^k). \tag{19}$$

Then the feasible set (13) *is reachable by solving a finite number of minimization problems in X with ε^k fixed for each problem.*

In this theorem, Garanzha et al. not only proved that there exists a regularization parameter sequence $\{\varepsilon^k\}_{k=0}^K$ leading to $F(X^K, 0) < +\infty$, but also provided an actual update rule for ε^k, refer to [Garanzha et al. 2021, Eq. (6)]. Inspired by these results, we formulate a very similar theorem allowing us to build maps with bounded distortion in a finite number of steps.

We also provide a way to build an increasing sequence $\{t^k\}$ to be used in Alg. 2–line 3: denote by $f_i(X^{k+1})$ the distortion for the element i, and by f_+^{k+1} the maximal distortion value over the mesh X^{k+1}, $f_+^{k+1} := \max_i f_i(X^{k+1})$. We propose to use the following update rule for t^{k+1}:

$$t^{k+1} := t^k + \sigma \frac{1 - t^k f_+^{k+1}}{f_+^{k+1}}, \tag{20}$$

where $0 < \sigma < 1$ is again the performance index of the minimization toolbox. Clearly, formula (20) does not involve t^*. Alg. 2 along with this update rule define our quasi-isometric stiffening (QIS) algorithm.

Now we are ready to formulate the stiffening theorem.

Theorem 2. *Let us suppose that the feasible set of bounded distortion meshes* (17) *is not empty, namely there exists a constant $0 < t^* < 1$ and a mesh X^* satisfying $W(X^*, t^*) < +\infty$. We also suppose that for solving $X^{k+1} \leftarrow \arg\min_X W(X, t^k)$ we have a minimization algorithm satisfying the following efficiency conditions for some $0 < \sigma < 1$:*
For each iteration k,

- ***either*** *the essential descent condition holds*

$$W(X^{k+1}, t^k) \le (1 - \sigma) W(X^k, t^k), \tag{21}$$

- ***or*** *the vector X^k satisfies the quasi-minimality condition:*

$$\min_X W(X, t^k) > (1 - \sigma) W(X^k, t^k). \tag{22}$$

Then the feasible set (17) *is reachable by solving a finite number of minimization problems in X with t^k fixed for each problem. In other words, under a proper choice of the continuation parameter sequence $\{t^k\}_{k=0}^K$, we obtain $W(X^K, t^*) < +\infty$.*

Proof. The main idea is very simple: update rule (20) defines an increasing sequence $\{t^k\}_{k=0}^\infty$. We will show that the corresponding sequence $\{W(X^k, t^k)\}_{k=0}^\infty$ is bounded from above. Then we can prove that the admissible set (17) is reachable in a finite number of steps by a simple *reductio ad absurdum* argument.

More precisely, if the feasible set is not reachable, then $W(X^k, t^k)$ must grow without bounds, which contradicts the boundedness. To prove that $\{W(X^k, t^k)\}_{k=0}^\infty$ is bounded from above, we analyze the behavior at some iteration k. First of all, if we couple update rule (20) with the fact that for any $t_2 > t_1$ the ratio $\frac{1 - t_1\psi}{1 - t_2\psi}$ is an increasing function of the argument ψ, we can see that the following inequality holds:

$$(1 - \sigma) W(X^{k+1}, t^{k+1}) \le W(X^{k+1}, t^k). \tag{23}$$

More precisely,

$$W(X^{k+1}, t^{k+1}) = \sum_i \frac{1 - t^k f_i(X^{k+1})}{1 - t^{k+1} f_i(X^{k+1})} \frac{f_i(X^{k+1})}{1 - t^k f_i(X^{k+1})} \operatorname{vol}(T_i) \leq$$

$$\sum_i \frac{1 - t^k f_+(X^{k+1})}{1 - t^{k+1} f_+(X^{k+1})} \frac{f_i(X^{k+1})}{1 - t^k f_i(X^{k+1})} \operatorname{vol}(T_i) = \frac{1}{1-\sigma} W(X^{k+1}, t^k).$$

Then, at each iteration k, either condition (21) or condition (22) must be satisfied. Let us consider both cases.

Cond. (21) **holds:** in this case function W actually decreases. Eq. (23) combined with Cond. (21) directly imply that

$$W(X^{k+1}, t^{k+1}) \leq W(X^k, t^k).$$

Cond. (22) **holds:** under assumption that $t^{k+1} < t^*$, by combining Eq. (23) and Cond. (22), we obtain:

$$W(X^{k+1}, t^{k+1}) \leq \frac{1}{1-\sigma} W(X^{k+1}, t^k) \leq$$

$$\frac{1}{(1-\sigma)^2} \min_X W(X, t^k) \leq \frac{1}{(1-\sigma)^2} W(X^*, t^k) \leq$$

$$\leq \frac{1}{(1-\sigma)^2} W(X^*, t^*)$$

To sum up, in the first case the value of function W decreases, and in the second case it does not exceed a global bound. Therefore, the sequence $\{W(X^k, t^k)\}_{k=0}^{\infty}$ is bounded from above.

Now we are ready for the main result. Suppose that for an infinite sequence $\{X^k, t^k\}_{k=0}^{\infty}$ we never reach the given quality bound t^*. In other words, we have $t^k < t^*,\ k = 0, \dots \infty$. Then the following inequality holds (apply formula (20) k times):

$$t^k - t^0 = \sigma \sum_{j=0}^{k-1} \left(\frac{1}{f_+^{j+1}} - t^j \right) \leq 1.$$

In the infinite sum, each term is strictly positive, hence we can extract a subsequence:

$$1 - t^{j_s} f_+^{j_s+1} \to 0^+.$$

This fact obviously contradicts the boundedness of the functional, allowing us to conclude our proof. □

Remark 1. *An important corollary of Theorem 2 is that, provided that the admissible set* (17) *is not empty, there exists an iteration $K < \infty$ such that the global minimum of the function $W(X, t^K)$ belongs to the admissible set. The proof is rather obvious: suppose we have an idealized minimizer such that $X^{k+1} = \arg\min_X W(X, t^k)$. This minimizer always satisfies the conditions of Theorem 2, therefore it can reach the distortion bound in a finite number of steps.*

N.B. While a solver that satisfies the optimization efficiency condition for *arbitrary* configurations does not exist, in practice, for the lowest distortion map problem, we did not encounter a situation where the update rule (20) leads to a solution worse than we were able to build with other methods. So, having a usual optimization toolbox (like Newton or L-BFGS), we need to estimate its efficiency σ, which is not known in advance. Garanzha et al. [2021] suggest computing the local descent coefficient for each minimization step, and so do we. Instead of σ in the update rule (20), we use σ^k defined as follows:

$$\sigma^k := \max\left(1 - \frac{W(X^{k+1}, t^k)}{W(X^k, t^k)}, \sigma_0\right),$$

where $\sigma_0 > 0$ is a constant (we chose $\sigma_0 = 1/10$ in all our experiments).

Characterization of the stiffness of variational problems is a challenging task. We have managed, however, to demonstrate that the positive definite part of the Hessian matrix of functional (16) is spectrally equivalent to the finite element approximation of the Dirichlet functional on the same computational mesh X, provided that the function $W(X, t)$ is bounded from above. More formally, let us denote the Hessian matrix as $H = \frac{\partial^2 W}{\partial X \, \partial X^\top}$. Then, since function W is the sum of elementary terms depending on k-th cell distortion $f(J_k)$, and distortion measure f is polyconvex, it can be written as

$$f(J_k) = \Phi(J_k, D_k), \quad D_k = \det J_k.$$

Finally, we define the positive definite part H^+ of the Hessian matrix H by neglecting second derivatives of D_k with respect to X when implementing chain rule for second derivative of f. Details of derivations are provided in the supplementary materials. We prove that the following statement holds:

Theorem 3. *Suppose that $W(X, t) < C$, where $0 < \vartheta < 1$, $0 < t < 1$, $C \geq 1$ are prescribed constants . Then there exist constants $k_1 > 0, k_2 > 0$, depending only on ϑ, t, C, such that*

$$k_1 \mathcal{D}_h(X) < \frac{1}{2} X^\top H^+ X < k_2 \mathcal{D}_h(X),$$

where $\mathcal{D}_h(X)$ is the discrete Dirichlet functional for standard simplicial linear finite elements which approximates Dirichlet functional

$$\mathcal{D}(x(\xi)) = \frac{1}{2} \int_\Omega \sum_i |\nabla x_i|^2 \, d\xi.$$

This theorem demonstrates stability of the positive definite part of the Hessian matrix when compressing admissible set, refer to the supplementary materials for a detailed proof.

5 VALIDATION

This section presents a thorough testing of our method. First of all, we show that QIS provides the best guess for equidistributed distortion. We start from the ground truth flattening of a hemisphere (§5.1), and then run a stress test for the mesh independence study of equidistribution of distortion (§5.2). Finally, we perform a massive testing of QIS against five competing methods over a large dataset (§5.3).

5.1 Flat Earth mathematics

In this section, we apply the equidistribution principle to hyperelastic deformations to tackle the classical cartography problem: the construction of a scale-optimal map of the unit hemisphere. The main goal is to establish a ground truth to assess QIS behavior.

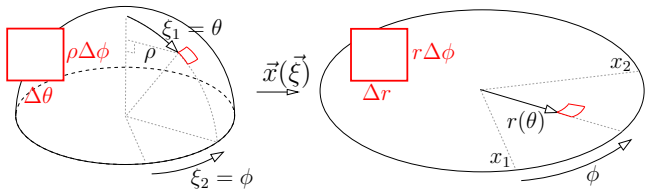

Fig. 3. The optimal flattening of the hemisphere is rotation-invariant, so we can represent it by a function $r(\theta)$, where r is the polar radius in the disc and θ is the colatitude on the hemisphere. Red squares illustrate the corresponding area differentials.

5.1.1 Ground truth. Consider a unit radius northern hemisphere H with colatitude $0 \le \theta \le \frac{\pi}{2}$ and longitude $0 \le \phi < 2\pi$. Due to the symmetry of the problem, H is to be mapped to a disc D. We denote by r, ϕ the polar coordinates in the disc. The mapping is rotation-invariant, so ϕ is irrelevant, and the mapping can be represented by a function $r(\theta)$. Refer to Fig. 3 for an illustration.

The azimuthal equidistant projection, based on the isometric flattening of meridians, has been known since ancient Egypt. However, it was only in [1969] that Milnor proved that it provides the best max-to-min scale ratio. In our notation, this map can be written as $r(\theta) = \theta$. The scale ranges from 1 to $\pi/2$, hence after normalizing this mapping by $\sqrt{2/\pi}$, we obtain the optimal (in the sense of (1)) mapping

$$r(\theta) = \sqrt{\frac{2}{\pi}}\theta. \tag{24}$$

with a quasi-isometry constant $K = \sqrt{\pi/2}$.

Recall that in this work, we consider hyperelastic deformations as a proxy problem; let us build the reference deformation by applying the equidistribution principle, see how it compares to Eq. (24), and get a ground truth for QIS.

The area differential on the hemisphere is defined by

$$d\sigma = \sin\theta\, d\theta\, d\phi,$$

and for a general rotation-invariant map $r(\theta)$, $\dot{r} > 0$, the singular values of the Jacobian matrix J of the mapping $H \to D$ are equal to

$$\sigma_1 = \dot{r}, \quad \sigma_2 = \frac{r}{\sin\theta}.$$

Therefore

$$\det J = \frac{\dot{r}r}{\sin\theta}, \qquad \operatorname{tr} J^\top J = \dot{r}^2 + \frac{r^2}{\sin^2\theta},$$

and the distortion measure (6) with $\vartheta = 1/2$ can be written as a function of r and θ:

$$f(J) = \frac{1}{4}\left(\frac{r}{\dot{r}\sin\theta} + \frac{\dot{r}\sin\theta}{r}\right) + \frac{1}{4}\left(\frac{\dot{r}r}{\sin\theta} + \frac{\sin\theta}{\dot{r}r}\right). \tag{25}$$

The equidistribution principle requires finding a function $r(\theta)$ such that the distortion f is constant for all values of the parameter θ. In other words, we want to solve the differential equation

$$f(\dot{r}(\theta), r(\theta)) = c \qquad \forall\theta \in [0, \pi/2] \tag{26}$$

for some constant c. We have numerically found the lowest possible value of c, thus obtaining the ground truth for QIS:

$$c = 1.030465 \pm 10^{-6}. \tag{27}$$

To compute (27), first we minimized the mean distortion by solving

$$\arg\min_{r(\theta)} \int_0^{\pi/2} f(J)\sin\theta\, d\theta. \tag{28}$$

The maximum distortion over the solution gave us the starting point, and then we used the shooting method to solve (26) and match the boundary conditions on the right, where the free parameter is the constant c. It is done by minimizing c to avoid the appearance of singularities inside the computational domain. Fig. 4 shows the deviation of the ground truth solution from (24).

It is easy to see that our proxy problem is very close to the best quasi-isometry Prob. (1). The source of discrepancy is the free boundary condition for elastic material that results in nonlinear Neumann-type boundary conditions on the free boundary. This boundary condition is expressed as $\dot{r}(\pi/2) = 1$, which is not compatible with the optimal solution (24), refer to supplementary materials for more details.

5.1.2 How well does QIS behave? Having found the ground truth (27), it is time to flatten surface meshes. We have built a regular triangulation of the hemisphere (so-called "uv-mesh" made of 10k triangles), and then flattened it with 4 different methods: simplex assembly [Fu and Liu 2016]), elastic deformation (ED, Prob. (2)), advanced MIPS [Fu et al. 2015]), and QIS. The results are summarized in Fig. 5. Even if all four methods optimize for exactly the same notion of the distortion measure (6), the performance is very different. The optimization scheme of simplex assembly (SA) results in a very noisy solution that is very different from the mean distortion minimizer (ED).

QIS creates an almost constant distortion field with a maximum distortion equal to 1.0307, differing from the ground truth (27) in the fifth digit. Advanced MIPS, however, fails to reach such quality. Recall that the idea of AMIPS is to penalize large values of distortion with the exponential law by minimizing

$$\int \exp(s\, f(J)),$$

where $s > 0$ is the stiffening parameter. Clearly, higher values of the constant s result in less maximum distortion. The authors do not provide a strategy of choice for the constant and propose a simple rule of thumb $s = 5$. We did our best to push s as high as possible: the solution in Fig. 5 was obtained with $s = 300$, we failed to go any higher. We had to give the QIS result as the initialization for AMIPS, and even then the optimizer had to handle 10^{150} energy values, something very delicate for double floating-point precision. Even with $s = 300$, AMIPS still lags behind in the sense that constant distortion is not attained and the maximum distortion is equal to 1.0377. It is important to note that this is a toy problem in controlled settings and we do not know how to attain large values of s in real-world scenarios.

5.2 Stress Test

As we have already mentioned, our approach is a discretization of a well-posed variational scheme, and it has an advantage that the type, size, and quality of mesh elements in the deformed object have

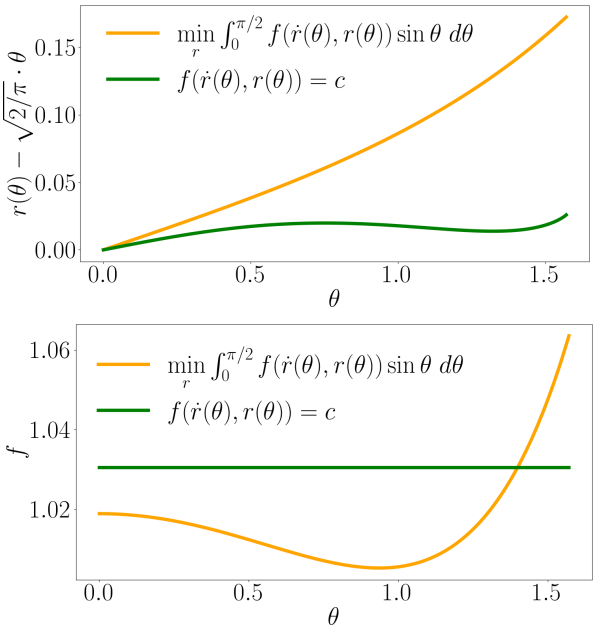

Fig. 4. Orange and green: two flattenings of the unit hemisphere. One minimizes the mean distortion and the other one minimizes the maximum distortion, thus satisfying the equidistribution principle. Top row: the deviation of the flattenings from the azimuthal equidistant projection, bottom row: corresponding distortion along a meridian.

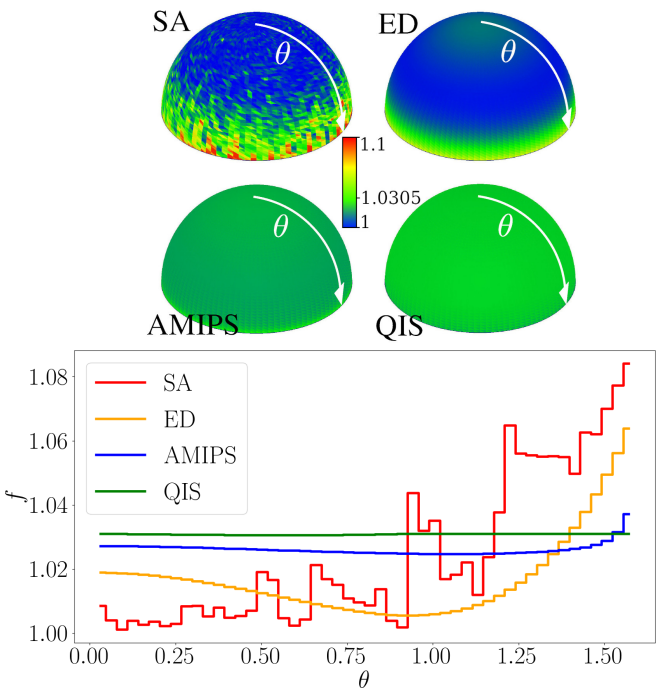

Fig. 5. Four flattenings over a regular 10k triangles uv-mesh of the unit hemisphere. Top: plot of the distortion fields. Bottom: the distortion field section along the zero meridian.

a weak influence on the computed deformation. To illustrate this point, we have computed free-boundary quasi-isometric maps on two different meshes of the same object. One mesh is a very regular triangulation with 73k facets, while the other one is of very bad quality (4k triangles, maximum triangle aspect ratio of 10^8). We start with elastic deformation (ED), see Fig. 6(a), (d). As expected, sharp peaks in the distortion field are present, which sharply grow with mesh refinement. Fig. 6(b), (e) show results for the QIS algorithm. The results are very similar; in both cases, QIS created plateaux. The height difference between cases (e) and (b) is about 10%, which is a very good result for an ugly triangulation (e). The material corresponding to plateaux behaves as a rigid one.

Since the essence of the QIS algorithm is the contraction of the set of admissible elastic deformations, the natural question is about the limit of contraction. Our experimental evidence suggests that when the first major plateau in the distortion measure distribution is reached, we get locked spots in the elastic material. The material around is still elastic, meaning that the admissible set is still not collapsed into a point, and the current solution is situated almost on the boundary of this set. Potentially, one can try to build a distribution of the parameter t in Ω, consisting of sets of plateaux with smooth transitions in order to make all cells of elastic material rigid, meaning that the admissible set is at last collapsed into a point.

A very similar idea was expressed in [Levi and Zorin 2014, §3]. The authors noticed that max-norm minimizers are generally not unique, and they proposed to fix a set of triangles for which the maximal distortion is reached, and then solve another max-norm optimization on the remaining triangles. Continuing this process, they eventually fix bounds for all triangles, decomposing them into groups with the same distortion.

Fig. 6(c) shows a prototype of this algorithm. Having obtained Fig. 6(b), we selected mesh cells inside and in the small vicinity of principal plateaux, and ran a bi-material stiffening, meaning that we had two different values of the parameter t depending on the material. This allowed contracting more for the remaining mesh, while locked spots behave like rigid inclusions that can move or rotate in the process of optimization. Evidently, allowing jumps in the parameter t results in non-smooth deformations, which is undesirable. The transition model for the parameter t between consecutive plateaux should be built allowing uniform locking of the material. It seems that the target least distortion solution means a locked state everywhere in the domain.

5.3 Massive Testing

To assess QIS behavior, we have performed a series of tests. We ran competing methods on a benchmark database. This section is structured as follows: for each adversary method, we closely inspect an example of two competing maps, and we provide statistics on several thousand other maps.

Recall that in our mapping pipeline, first we compute a foldover-free map with [Garanzha et al. 2021], and then we call QIS to

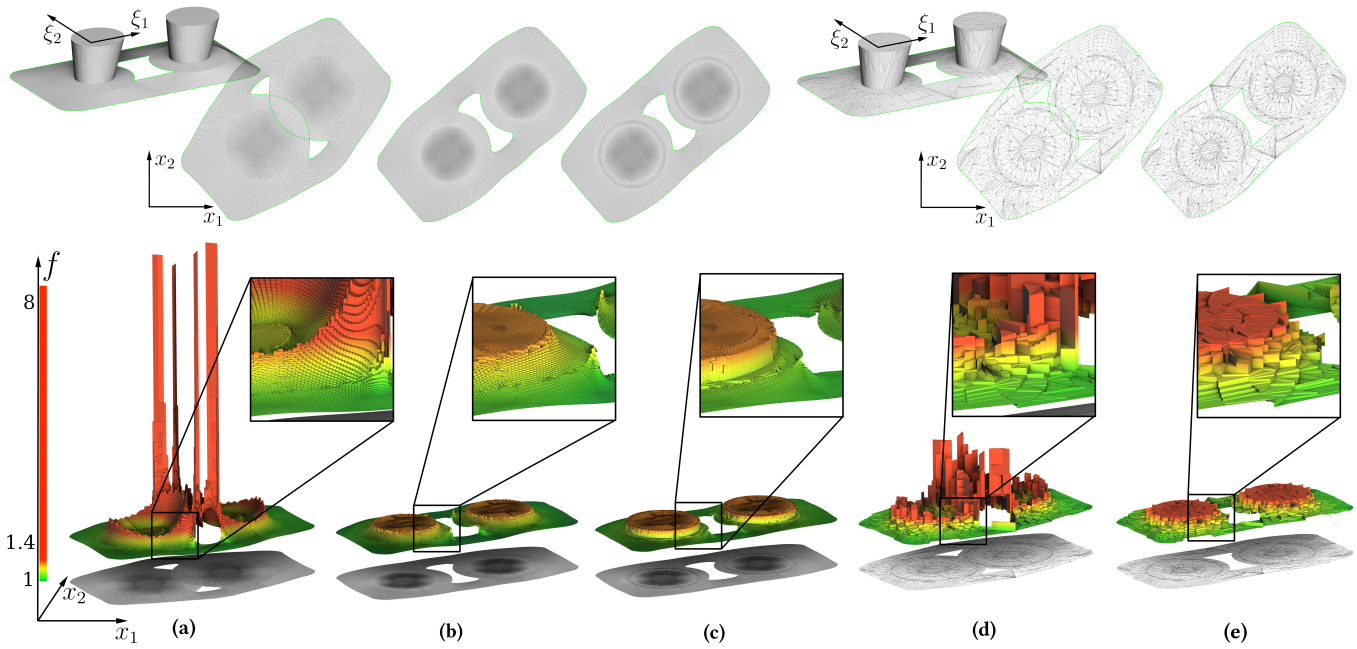

Fig. 6. Our method shows excellent stability irrespective of input mesh quality. Here we show five free-boundary flattenings for two input meshes: one has a very regular triangulation (73k triangles), while the other one has a very bad quality triangulation (4k triangles, maximum triangle aspect ratio of 10^8). Top row: flattenings, bottom row: plots of the distortion measure. (a) and (d) columns show elastic flattening results, (b) and (e) quasi-isometric stiffening, (c) bi-material stiffening. QIS suppresses peaks of distortion, replacing them by minimum height plateaux. Bi-material stiffening allows creating a second-order plateau.

compute a map with the lowest distortion. Fig. 1 and 7 provide a 2D and 3D example of quasi-isometric maps, both computed with $\vartheta = \frac{1}{2}$. It is easy to see that elastic deformations sacrifice the quality of several elements to minimize distortion on average, while QIS distributes the distortion evenly over the entire domain, thus effectively reducing the maximum distortion.

We ran the same procedure over the database provided by Du et al. [2020]. The benchmark contains 10743 triangulated meshes to flatten and 904 constrained boundary 3D mapping challenges. We have flattened the surfaces without imposing boundary constraints.

Fig. 8 summarizes the results: it shows the scatter plot for every mapping challenge from the database. In the top row, we plot the worst element quality given by the elastic deformation against the worst element quality in maps computed by QIS. It is easy to see that the plots are upper-triangular: without surprise, QIS is consistently better. Note the dots on the main diagonal in 3D mapping; they correspond to the deformations whose worst quality tetrahedron has all 4 vertices locked, leaving no degrees of freedom for improvement.

The bottom row of Fig. 8 provides a $\log - \log$ scatter plot of our running time versus mesh size for all the challenges from the database: for each run, the time varies from a fraction of a second to several minutes for the largest meshes. These times were obtained with a 12-core i7-6800K CPU @ 3.40 GHz.

We have also compared QIS maps against 5 competing methods:

(1) Simplex assembly [Fu and Liu 2016]
(2) Power law enhancement [Fang et al. 2021; Garanzha and Kudryavtseva 2019]
(3) Exponential law enhancement (SLIM+AMIPS [Fu et al. 2015])
(4) Large-scale bounded distortion mappings [Kovalsky et al. 2015]
(5) Strict minimizers for geometric optimization [Levi and Zorin 2014].

The comparison is made on free-boundary flattening of 10743 triangular meshes from the [Du et al. 2020] database. With two exceptions, we gave all methods the same input, i.e. fold-free maps provided by [Garanzha et al. 2021]. To avoid floating point errors, we had to give the QIS result as initialization for SLIM+AMIPS (more on that below), and simplex assembly does not require an initialization.

Simplex Assembly. First, we compare QIS to Simplex Assembly (SA) [Fu and Liu 2016]. Simplex assembly is a method to compute inversion-free mappings with bounded distortion on simplicial meshes. The idea is to project each simplex into the inversion-free and distortion-bounded space. After disassembling the mesh, the simplices are then assembled by minimizing the mapping distortion while keeping the mapping feasible. Fig. 9 provides a quality comparison of SA with our quasi-isometric ($\vartheta = \frac{1}{2}$) map for a free-boundary mapping of the "Lucy" mesh. This comparison is interesting for two reasons: first, SA offers an explicit optimization for the distortion

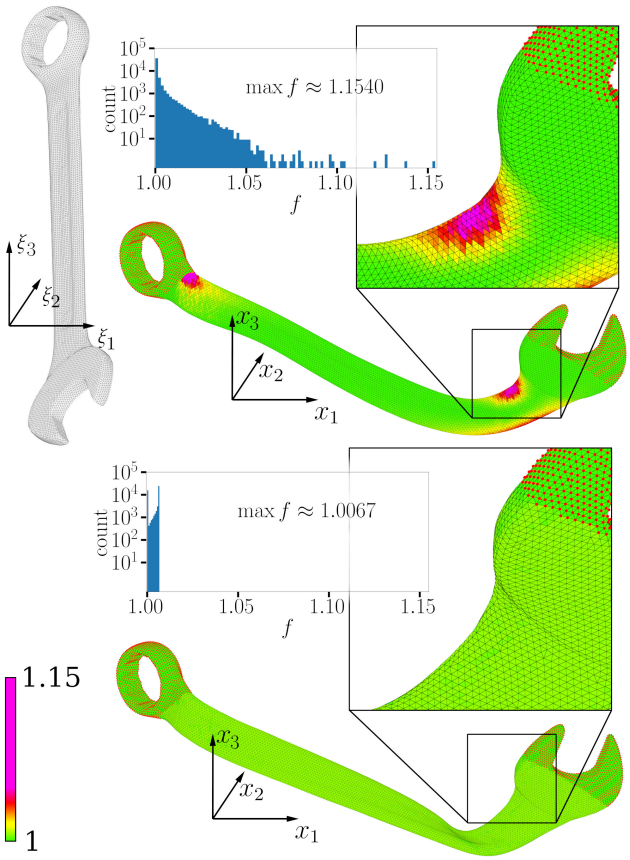

Fig. 7. Tetrahedral mesh deformation, locked vertices are shown in red. Top left: rest shape, top right: elastic deformation [Garanzha et al. 2021], bottom: QIS. Both methods optimize for the same distortion measure f, $\vartheta = \frac{1}{2}$, whose histograms are provided in the corresponding plots.

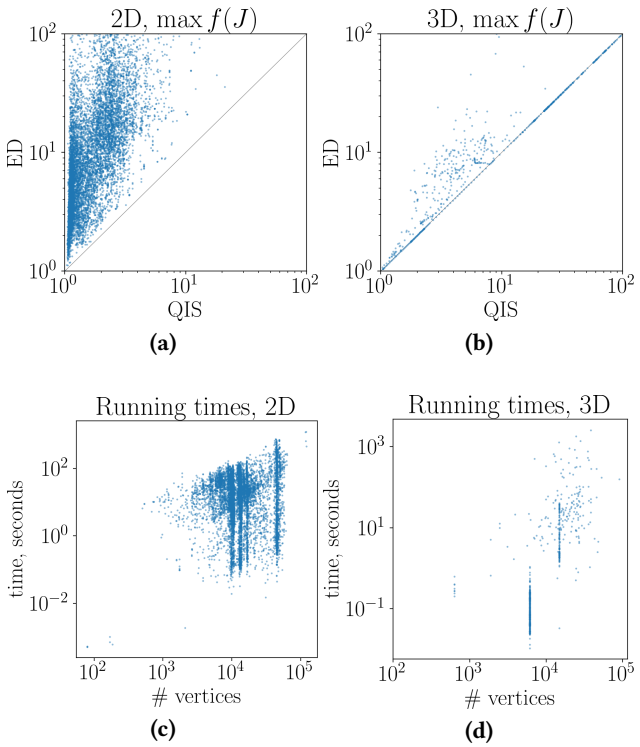

Fig. 8. Elastic deformation [Garanzha et al. 2021] vs QIS tested on 10743 triangular meshes (a) and 904 tetrahedral meshes (b) in the database provided by [Du et al. 2020]. The boundary is set free for 2D flattenings and locked for 3D. Top row: worst element quality comparison. Bottom row: QIS running times.

bound, and second, in 2D, Fu et al. use exactly the same distortion measure f as we do, allowing us to directly compare the distortion.

Fig. 10-a shows the quality plot of SA maps against QIS. Without any exceptions, QIS produces better results. Note that SA ignores the locally-injective initializations we have offered; it failed to produce valid maps in 15 cases out of 10743.

Power Law Enhancement. Fang et al. [2021] attempt to improve worst-element distortions by formulating a regularized min-max optimization for IDP by applying a p-norm extension to the symmetric Dirichlet (SD) energy with an exponential factor $p > 1$. Garanzha et al. [2019] used the same idea, but for the distortion measure f instead of SD.

We ran the IDP code on one 3D mesh from their examples and implemented power law enhancement for 2D flattening. In 3D, we deformed a cylinder tetrahedral mesh. We applied two bone handles (two thin boxes of interior axis vertices) to bend it. Fig. 11 shows the comparison of our results with IDP. Locked vertices are shown in red.

As advised by Fang et al., we chose $p = 5$. It improves slightly the worst-element distortion. As shown in Fig. 11–left, the stress is concentrated around the locked vertices (shown in magenta). Our optimization ($\vartheta = \frac{1}{2}$) allows us to dissipate the stress over a larger area, thus improving both distortion measures: the maximum stretch decreases from 5.05 to 1.94, and the minimum scale increases from 0.36 to 0.72.

For surface flattening, we have minimized $\int f^5$, which allows us to directly compare the maximum values of f over the maps. Fig. 10-b shows the quality plot of power law enhancement maps against QIS. Once again, the plot is upper triangular, where QIS produces better results for every challenge.

Exponential law enhancement. For our next step, we have implemented the SLIM+AMIPS combo [Fu et al. 2015; Rabinovich et al. 2017][4]. The idea is to minimize $\int \exp(s \cdot f(J))$, thus penalizing large distortion values by the exponential law. We have already provided an example of AMIPS flattening in §5.1, where in very controlled settings we have managed to use $s = 300$. For the database, we have used $s = 5$, and the results can be seen in Fig. 10-c. We have used the default trade-off parameter $\vartheta = \frac{1}{2}$. The results are better than

[4]Note a typographic error in [Rabinovich et al. 2017, Table I]: AMIPS optimizes $\exp(s \cdot f(J)) = \exp\left(\frac{s}{4}\left(\frac{\sigma_1}{\sigma_2} + \frac{\sigma_2}{\sigma_1} + \sigma_1\sigma_2 + \frac{1}{\sigma_1\sigma_2}\right)\right)$.

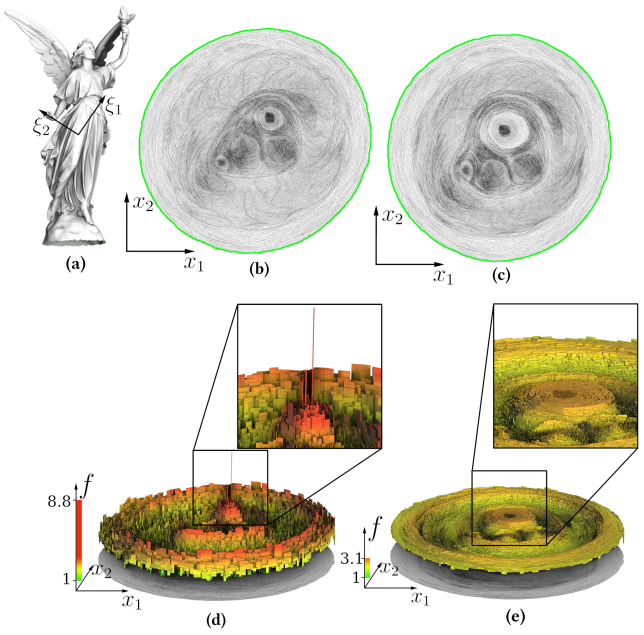

Fig. 9. Two quasi-isometric maps for the "Lucy" mesh (a): simplex assembly (b) vs QIS (c). Both methods optimize for the same distortion measure f, $\vartheta = \frac{1}{2}$, whose plot is given in (d) and (e).

for the power law enhancement, but still not as good as QIS. Again, the plot is strictly upper triangular. Note that we had to use the QIS result as the initialization to avoid failures. Even for a quite low value of s, due to floating point error problems, the SLIM+AMIPS combo fails to produce valid maps in 1796 cases out of 10743 if foldover-free maps [Garanzha et al. 2021] are used as input.

Large-scale Bounded Distortion Mappings. Our next test is LBD [Kovalsky et al. 2015]. Given an input map (potentially with folds), LBD looks for an injective map as close as possible to the input map, but satisfying some constraints such as the orientation as well as distortion bounds. Generally speaking, the problem of minimizing an energy subject to bounded distortion constraints is known to be difficult and computationally demanding. LBD alternates between energy minimization steps and projection to the constraints.

For an illustration (refer to Fig. 12), we took the example from the source code provided by the authors. The 3D surface to flatten is a regular simplicial mesh of a rectangular patch that was lifted and noised. Since LBD has an explicit optimization of the distortion bounds, the comparison is of particular interest. Note that a) LBD needs a user-provided bound and b) it optimizes for the Jacobian condition number. Therefore, we have computed QIS with $\vartheta = 0$, computed the worst condition number $\max \frac{\sigma_1}{\sigma_2} = 1.4252$, and ran LBD with that bound. The result is very interesting: up to numerical errors, LBD managed to find a flattening with the prescribed bound, however, the vast majority of the triangles are at the bound, whereas QIS has globally very low distortion and only a few bad triangles. Clearly, to have the best map, not only do we need to reach the feasible set (17), but also find its central point. Apparently, elastic

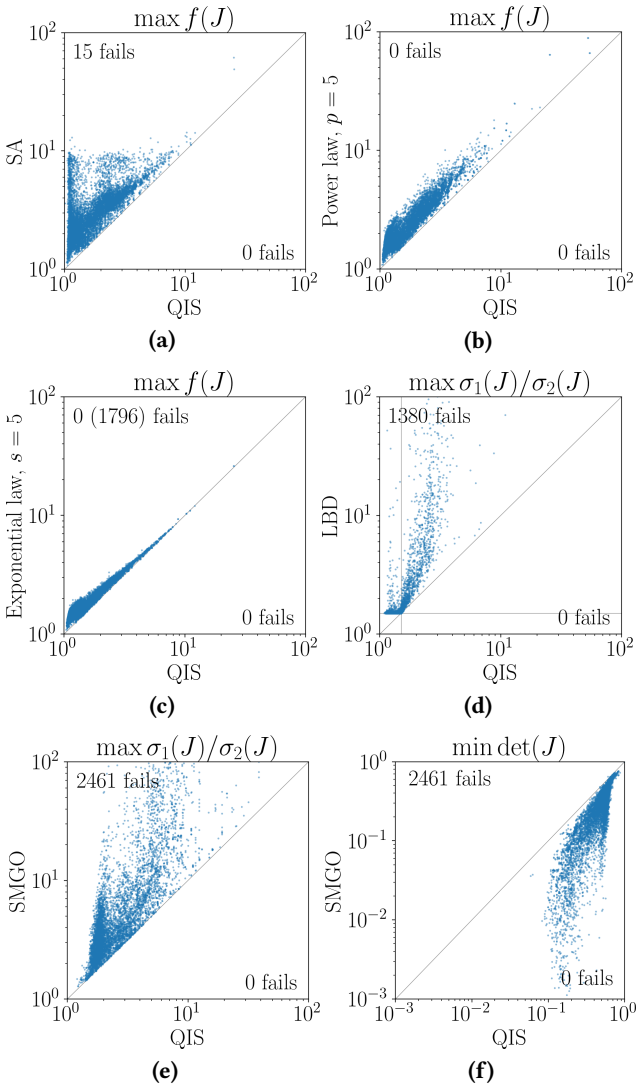

Fig. 10. Quality of QIS free-boundary flattening against 5 competing methods. The comparison is made on 10743 triangular meshes from the [Du et al. 2020] database. (a) Simplex assembly [Fu and Liu 2016] vs QIS; (b) power law enhancement [Fang et al. 2021] vs QIS; (c) SLIM+AMIPS [Fu et al. 2015] vs QIS; (d) Large-scale bounded distortion mappings [Kovalsky et al. 2015] vs QIS. The vertical and horizontal lines in (e) highlight the Jacobian condition matrix bound 1.5 chosen by the authors of LBD. (e) and (f) Strict minimizers for geometric optimization [Levi and Zorin 2014] vs QIS.

deformations provide a better alternative to ad-hoc projections to the highly non-convex set of max distortion constraints.

Fig. 10-d provides a comparison of LBD vs QIS over the database. Since LBD optimizes for the Jacobian condition number, we had to compute QIS with $\vartheta = 0$ for all maps. For LBD, we used the default bound $\max \frac{\sigma_1}{\sigma_2} = 1.5$. The plot shows a direct comparison of corresponding condition numbers, and it is strictly upper triangular,

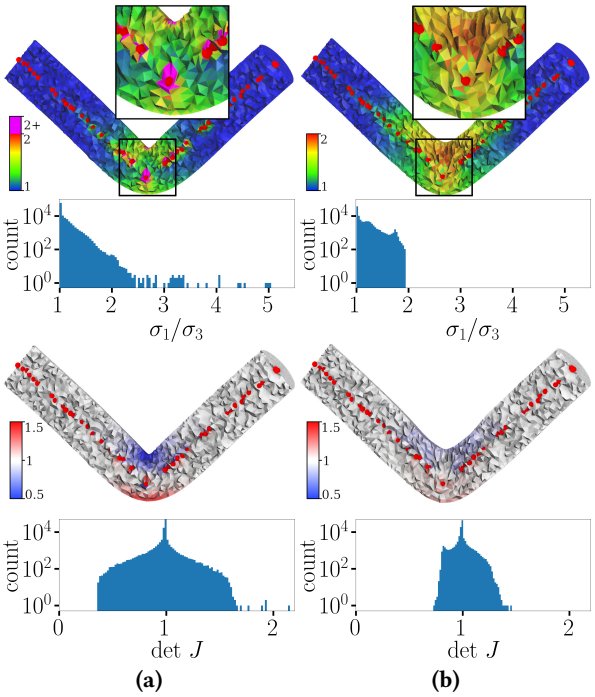

Fig. 11. Bending test for a tetrahedral mesh of a cylinder, locked vertices are shown in red. (a): IDP [Fang et al. 2021], (b): QIS deformation with $\theta = \frac{1}{2}$. From top to bottom: Jacobian matrix condition number and the Jacobian determinant are shown in histograms and corresponding color plots.

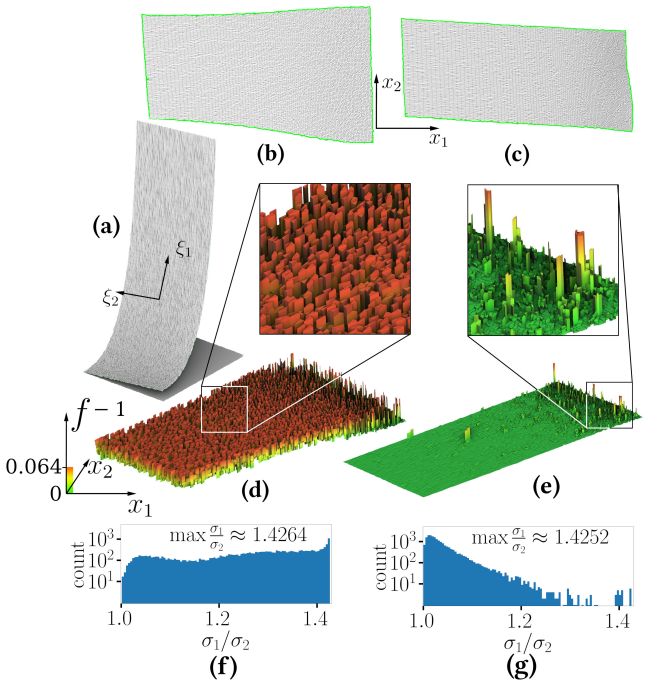

Fig. 12. Comparison of LBD [Kovalsky et al. 2015] vs QIS. (a): The 3D surface to flatten is a regular triangular mesh of a square patch that was lifted and noised. The distortion bound for LBD is set by the QIS result. (b): The map obtained by LBD. (c): QIS ($\vartheta = 0$). (d) and (e): Plots of the distortion measure. (f) and (g): Jacobian matrix conditioning histograms.

so QIS is consistently better. Note that in 1380 cases LBD failed to find a map with the prescribed bound, while QIS succeeded. In addition to that, despite the fact that LBD was provided with locally injective maps, it returned invalid (with folds) maps in 5412 cases out of 10743. Note that these failures are not entirely due to the unreachability of the deformation bound: QIS has found 66 valid maps respecting the bound out of 5412 rejected by LBD.

For the sake of completeness, we have also tried to run LBD with a *lower* bound than the one found by QIS (we lowered it by 1%). LBD did not manage to find a single map respecting this bound.

Strict Minimizers For Geometric Optimization. Our final competitor is SMGO [Levi and Zorin 2014]. Given an input triangulated surface and a set of reference rotation matrices $\{R_i\}$, SMGO computes a flattening by solving the following max-norm optimization problem:

$$\min k, \quad \text{subject to} \quad \|J_i - R_i\|_F \le k \quad \forall i \in 1 \dots \#T.$$

This is a second-order cone programming problem, and it can be solved with MOSEK [Andersen and Andersen 2000].

As with other competitors, we ran it over the entire database. To find the target rotation matrices, we took the foldover-free maps and extracted the rotations from the Jacobi matrices by solving the orthogonal Procrustes problem. We chose to plot the maximum Jacobian condition number and the minimum Jacobian determinant. Fig. 10-e and 10-f summarize the results. Since SMGO and QIS optimize for different distortion measures, we did not expect strictly triangular plots. Indeed, SMGO has found two maps with better stretch and one map with better scale (but not both at the same time). It is to be noted that this plot was made with fixed target rotation matrices, and if we applied local-global ARAP iterations to update the target, the results would be significantly better for SMGO. However, we chose not to do so due to the very high computational cost involved in doing it over the database.

Note that SMGO has no guarantees of injectivity: despite the fact that the target rotations were extracted from foldover-free maps, SMGO failed to find a valid map in 2461 cases out of 10743. Fig. 13 illustrates the problem: a valid map can be found only if the input frame field has sufficiently low distortion.

6 CONCLUSION

We have formulated a set of variational problems that potentially cover the complete technological chain for constructing optimal mappings and deformations with fixed as well as free boundaries. We started with the continuation problem w.r.t. the parameter ε. This minimization allows us to compute optimal in the average deformations. We then formulated a continuation problem for the worst distortion measure minimization (quasi-isometric stiffening, QIS), which retains polyconvexity and smoothness of deformation. The stiffening technique tends to suppress peaks in the distortion

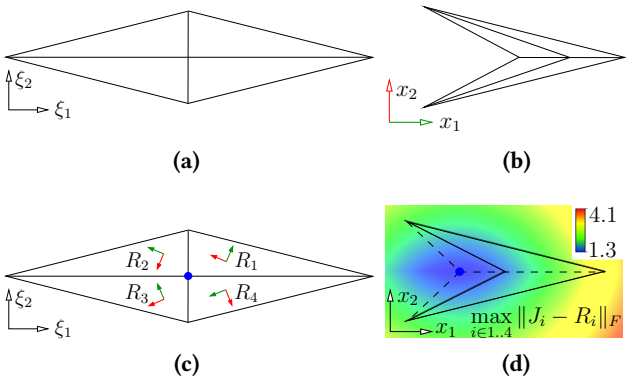

Fig. 13. Failure case for [Levi and Zorin 2014] for a constrained boundary mapping. Having a valid map from the reference surface (a) to a flattening (b), we can extract target rotation matrices $\{R_i\}_{i=1}^4$ shown in (c). Plot (d) shows the maximum distortion over the triangulation as a function of the blue vertex position. The minimum is attained outside of the target domain, implying foldovers in the output.

measure distribution in the computational domain and replaces them with plateaus of minimal height. Surprisingly, for some classical problems such as the optimal map on the hemisphere, our algorithm provides hyperelastic deformation with a constant distortion measure (deformation energy density). This is the first known example of an exact 2D implementation of the equidistribution principle in hyperelasticity. Throughout all stages, we took care to demonstrate that a finite number of basic optimization steps are sufficient to solve the problem, given certain efficiency assumptions on the optimization toolbox. We illustrated the performance of our algorithm with challenging 2D and 3D numerical tests.

ACKNOWLEDGMENTS

This work is supported by the Ministry of Science and Higher Education of the Russian Federation, project No 075-15-2020-799.

APPENDIX

A EQUIDISTRIBUTION PRINCIPLE AND SCALE-OPTIMAL MAPS

The construction of deformations that uniformly minimize deviation from isometry naturally leads to the idea that the optimal solution should, in some form, obey the equidistribution principle.

Usually, the formulation of the equidistribution principle is attributed to de Boor: in [1973], he computed the optimal location of spline control points for a max-norm minimization of interpolation error. Let us consider the simplest illustration of the idea. Suppose we have a uniformly bounded weight function $w(x)$

$$c_1 < w(x) < c_2, \qquad c_1 > 0, \int_0^1 w(x)\,dx = 1.$$

This function may play the role of an error indicator for rather abstract interpolation problem, as discussed in [1973].

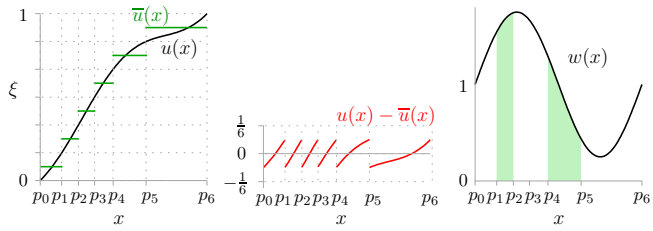

Fig. 14. Equidistribution principle in action: we place the nodes $p_0 \dots p_N$ in the way to obtain the optimal max-norm approximation of $u(x)$ by a piecewise constant function $\overline{u}(x)$ (left). The approximation error $|u(x)-\overline{u}(x)|$ does not exceed $\frac{1}{2N}$ (middle). The nodes are placed in the way to distribute the error indicator measure $w(x) = \dot{u}(x)$ evenly over all segments, i.e. $w(x)$ has the same integral over all segments $[p_i, p_{i+1}]$ (right).

We want to find the position of $N-1$ nodes $p_i > p_{i-1}$, $p_0 = 0$, $p_N = 1$ forming the following equidistribution mesh of the unit segment $[0,1]$:

$$\int_{p_i}^{p_{i+1}} w(x)\,dx = \frac{1}{N} \tag{29}$$

for all i, meaning that the mesh cell is smaller for larger values of the error indicator. The equidistribution principle provides a simple quasi-optimal solution to the highly complicated problem of max-norm error minimization [1973].

To the best of our knowledge, the first known application of equidistribution is Chebyshev alternance [1853]: the maximum deviation for the best max-norm polynomial approximation is a constant that is attained in a special set of Chebyshev points. A major advantage of resulting adaptive meshes is the ability to provide an optimal asymptotic error convergence rate for highly nonuniform and singular data, for example, corresponding to boundary layers in CFD [Bakhvalov 1969].

Now let us see how the equidistribution principle relates to the optimal mapping problem, refer to Fig. 14 for an illustration. Consider a mapping $x(\xi) : S \to S_x$, where S is the parametric domain $[0,1]$ endowed with the Cartesian metric tensor $G_\xi = 1$ (in 1D this is just a positive scalar function), and S_x is the image domain $[0,1]$ with the metric tensor $G_x = w^2$. Since we want $x(\xi)$ to preserve distances ($d\xi G_\xi d\xi = dx G_x dx$), it implies the following condition:

$$d\xi = w dx. \tag{30}$$

Integrating (30) over a parametric cell, we obtain the discrete isometry condition which coincides with (29). In 1D, the construction of the mapping is simple: consider the primitive function $u(x)$, $w(x) = \dot{u}(x)$. One can choose $u(0) = 0$, $u(x) = \xi(x)$, and $u(p_{i+1}) - u(p_i) = \frac{1}{N}$. The resulting mesh provides optimal piecewise-constant approximation of $u(x)$ by a function $\overline{u}(x)$ equal to $\frac{1}{2}(u(p_i) + u(p_{i+1}))$ on the interval $p_i < x < p_{i+1}$. The resulting approximation error function is a typical Chebyshev alternance [Tchebychev 1853]: the approximation error function has alternating sign, attaining its extremal values in the special set of points (Fig. 14–middle).

For the optimal mapping $x(\xi)$, a finite element approximation of the corresponding hyperelastic energy attains its absolute minimum and has a constant uniform distribution of the energy density over the parametric domain.

W. Huang [2001] provides a detailed discussion of the above equidistribution principle in multiple dimensions and its relation to deviation from the isometry relation (1). The idea that the equidistribution principle is equivalent to the construction of quasi-isometries between manifolds with metric was in the air, so it was recognized by many authors [Godunov et al. 1994; Huang 2001; Xu et al. 2011].

B CHOICE OF A PROXY

Our paper would not be complete without mentioning the fact that there is a large variety of distortion measures that can be used as a proxy for quasi-isometric mapping. The only condition a distortion measure has to comply with is the fact that the minimization of the distortion implies the minimization of the quasi-isometry constant.

B.1 Relationship between distortion measures and the quasi-isometry constant

More formally, given a deformation of a mesh, let us denote by K its quasi-isometry constant (maximal relative length distortion of the map). We want to estimate K through available values of $f(J_k)$, where J_k is the Jacobian matrix of k-the simplex of the mesh. First of all, we can write down following estimation, where $\sigma_j(J_k)$ stand for (ordered) singular values of J_k:

$$\frac{1}{K} \le \sigma_j(J_k) \le K. \tag{31}$$

In practice many different distortion measures can be used, but we are interested in those that satisfy $f(J_k) \ge 1$ and guarantee that inequality

$$f(J_k) < \frac{1}{t}, \quad t < 1 \tag{32}$$

implies (31). Naturally, for QIS scheme to work with a certain distortion measure f, we need the bound for K to tend to 1 as $t \to 1$.

In this section we analyze two possible choices of distortion measures, namely, Eq. (6) and Symmetric Dirichlet [Schüller et al. 2013].

Let us start with $f := (1-\vartheta)f_s + \vartheta f_v$. First of all, let us note that $f_s \ge 1$ and $f_v \ge 1$, so following inequalities hold:

$$\frac{(\frac{1}{d}\operatorname{tr} J_i^\top J_i)^{d/2}}{\det J_i} < c_1, \quad c_1 := \left(\frac{1-t\vartheta}{t(1-\vartheta)}\right)^{d/2}$$

and

$$\frac{1}{2}\left(\det J_i + \frac{1}{\det J_i}\right) < c_2, \quad c_2 := 1 + \frac{1-t}{t\vartheta}.$$

Reshetnyak's inequality [Reshetnyak 1966] implies that

$$\frac{\sigma_1(J_i)}{\sigma_d(J_i)} < c_1 + \sqrt{c_1^2 - 1}$$

From the above estimates we obtain the required bounds for $\sigma_j(J_i)$ [Garanzha 2000], so

$$K < \left(c_1 + \sqrt{c_1^2-1}\right)^{(d-1)/d}\left(c_2 + \sqrt{c_2^2-1}\right)^{1/d} \tag{33}$$

Note that in the 2D case and with $\vartheta = 1/2$ the bound for K takes the simplest form

$$K < \frac{(1+\sqrt{1-t})^2}{t}.$$

Indeed, Eq. (6) forces the quasi-isometry constant K to tend to unity when $t \to 1$, so our QIS technique is correct.

We use the measure (6), but QIS can be successfully applied to symmetric Dirichlet as well:

$$f_{SD}(J) := \|J\|_F^2 + \|J^{-1}\|_F^2.$$

Simple estimates for K can be also derived for f_{SD}. In terms of singular values, this measure can be written as follows:

$$f_{SD}(J) := \frac{1}{2d}\sum_{j=1}^{d}\left(\sigma_j^2 + \frac{1}{\sigma_j^2}\right).$$

In this case

$$K^2 < c_3 + \sqrt{c_3^2 - 1}, \quad c_3 = 1 + d\frac{1-t}{t}$$

The bound satisfies the requirements for QIS algorithm, so SD can be used in our stiffening scheme. Note, however, that even though SD distortion measure is a convex function of singular values, it is not clear whether it is polyconvex (convex with respect to minors of J).

B.2 Hemisphere flattening example

Let us compare behavior of f and f_{SD} on the ground truth example (§5.1). Note that Airy [1861] considered precisely the hemisphere flattening problem for different distortion measures. He provided detailed numerical tables comparing area distortion ("exaggeration"), shape distortion ("distortion") for equidistant azimuthal projection, stereographic projection, incompressible projection ("unchanged areas"), and BE (Balance of Errors) projection. For the practical construction of BE flattening of the hemisphere, he used Prob. (3) with simplified energy:

$$\begin{aligned} &\arg\min_{r(\theta)} \int_0^{\pi/2} f_{Airy}^{ARAP}(J)\sin\theta\, d\theta, \\ f_{Airy}^{ARAP}(J) &:= (\sigma_1 - 1)^2 + (\sigma_2 - 1)^2 \\ &= (\dot r - 1)^2 + \left(\frac{r}{\sin\theta} - 1\right)^2. \end{aligned} \tag{34}$$

Airy managed to find the solution of the Euler-Lagrange equations for (34) analytically:

$$r_A(\theta) = \sqrt{2}\tan\frac{\theta}{2} + 2\cot\frac{\theta}{2}\log\sec\frac{\theta}{2}. \tag{35}$$

The first term in (35) corresponds to the homogeneous solution of EL equations, we multiplied it by a $\sqrt{2}$ coefficient in order to satisfy Neumann boundary conditions on the outer boundary $\partial f_A/\partial \dot r = 0$ for $\theta = \pi/2$.

We computed six flattenings using three distortion measures (f, f_{SD}, and f_{Airy}^{ARAP}). Fig. 15 shows the maps that minimize the average distortion, and Fig. 16 shows the equidistributed distortion

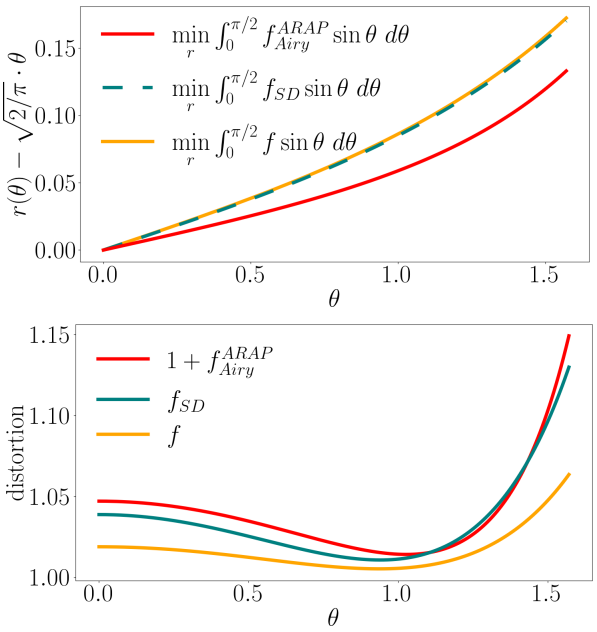

Fig. 15. Behavior of elastic hemisphere flattening under three different distortion measures. Here we minimize for the average distortion.

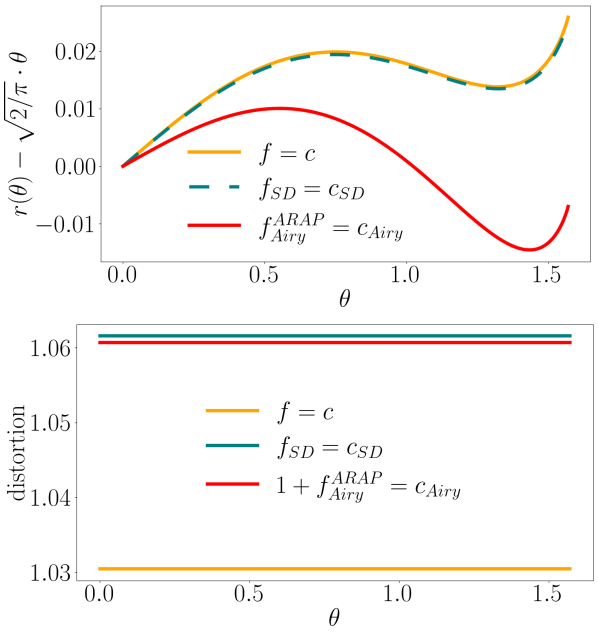

Fig. 16. Behavior of elastic hemisphere flattening under three different distortion measures. Here we minimize for the maximum distortion.

flattenings. In particular, we have found the following values for constants from (26):

$$c_{SD} = 1.061571 \pm 10^{-6},$$
$$c_{Airy} = 1.060678 \pm 10^{-6}.$$

All three distortion measures give pretty similar results, differing slightly from the equidistant azimuthal due to the aforementioned condition $\dot{r}(\pi/2) = 1$. Note that while QIS can use both f and f_{SD}, it cannot guarantee the absence of foldovers for ARAP-like energies due to the lack of barrier property. Fig. 13 provides an illustration. Note that in this test problem stiffening does not lead to sharp improvement of the scale error due to the fact that free boundary boundary conditions are responsible for largest part of contribution to the scale error.

REFERENCES

Noam Aigerman and Yaron Lipman. 2013. Injective and Bounded Distortion Mappings in 3D. *ACM Trans. Graph.* 32, 4, Article 106 (July 2013), 14 pages. https://doi.org/10.1145/2461912.2461931

Erling D. Andersen and Knud D. Andersen. 2000. *The Mosek Interior Point Optimizer for Linear Programming: An Implementation of the Homogeneous Algorithm.* Springer US, Boston, MA, 197–232. https://doi.org/10.1007/978-1-4757-3216-0_8

Nikolai Sergeevich Bakhvalov. 1969. The optimization of methods of solving boundary value problems with a boundary layer. *Ussr Computational Mathematics and Mathematical Physics* 9 (1969), 139–166.

John M Ball. 1976. Convexity conditions and existence theorems in nonlinear elasticity. *Archive for rational mechanics and Analysis* 63, 4 (1976), 337–403.

David Bommes, Henrik Zimmer, and Leif Kobbelt. 2009. Mixed-Integer Quadrangulation. *ACM Trans. Graph.* 28, 3, Article 77 (jul 2009), 10 pages. https://doi.org/10.1145/1531326.1531383

Mario Bonk and Urs Lang. 2003. Bi-Lipschitz parameterization of surfaces. *Math. Ann.* 327, 1 (2003), 135–169.

Pafnuty Lvovich Chebyshev. 1856. Sur la construction des cartes géographiques. *Bulletin de la classe physico-mathématique de l'Académie Impériale des sciences de Saint-Pétersbourg* VIV (1856), 257–261. Reprinted in P. L. Tchebychef, Œuvres I, Chelsea, New York, 1962, pp. 233-236 and 239-247.

Edward Chien, Zohar Levi, and Ofir Weber. 2016. Bounded Distortion Parametrization in the Space of Metrics. *ACM Trans. Graph.* 35, 6, Article 215 (nov 2016), 16 pages. https://doi.org/10.1145/2980179.2982426

P.G. Ciarlet and G. Geymonat. 1982. Sur les lois de comportement en elasticite non-lineaire compressible. *C.R. Acad.Sci. Paris Ser.II* 295 (1982), 423 – 426.

Philippe Ciarlet and Jindrich Necas. 1985. Unilateral problems in nonlinear, three-dimensional elasticity. *Arch. Rational Mech. Anal.* 87 (1985), 319–338. https://doi.org/10.1007/BF00250917

Philippe G. Ciarlet. 1988. *Mathematical Elasticity: Three-Dimensional Elasticity.* Number v. 1 in Studies in mathematics and its applications. North-Holland.

Carl de Boor. 1973. *Good Approximation by Splines with Variable Knots.* Birkhäuser Basel, Basel, 57–72. https://doi.org/10.1007/978-3-0348-5979-0_3

A. Dinchenko. 1938. Chebyshev projection for the Soviet Union (in Russian). *Geodesist* 10 (1938), 4–14. https://elib.rgo.ru/safe-view/123456789/222079/1/MDAwMDAzMDZfR2VvZGV6aXN0IOKEljEwLnBkZg==

Xingyi Du, Noam Aigerman, Qingnan Zhou, Shahar Z. Kovalsky, Yajie Yan, Danny M. Kaufman, and Tao Ju. 2020. Lifting Simplices to Find Injectivity. *ACM Trans. Graph.* 39, 4, Article 120 (July 2020), 17 pages. https://doi.org/10.1145/3386569.3392484

G. B. Airy Esq. 1861. LIII. Explanation of a projection by balance of errors for maps applying to a very large extent of the earth's surface; and comparison of this projection with other projections. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 22, 149 (1861), 409–421. https://doi.org/10.1080/14786446108643179

Yu Fang, Minchen Li, Chenfanfu Jiang, and Danny M. Kaufman. 2021. Guaranteed Globally Injective 3D Deformation Processing. *ACM Trans. Graph. (SIGGRAPH)* 40, 4, Article 75 (2021).

Xiao-Ming Fu and Yang Liu. 2016. Computing Inversion-Free Mappings by Simplex Assembly. *ACM Trans. Graph.* 35, 6, Article 216 (Nov. 2016), 12 pages. https://doi.org/10.1145/2980179.2980231

Xiao-Ming Fu, Yang Liu, and Baining Guo. 2015. Computing Locally Injective Mappings by Advanced MIPS. *ACM Trans. Graph.* 34, 4, Article 71 (jul 2015), 12 pages. https://doi.org/10.1145/2766938

Vladimir Garanzha. 2000. The barrier method for constructing quasi-isometric grids. *Computational Mathematics and Mathematical Physics* 40 (2000), 1617–1637.

Vladimir Garanzha and Igor Kaporin. 1999. Regularization of the barrier variational method of grid generation. *Comput. Math. Math. Phys.* 39, 9 (1999), 1426–1440.

Vladimir Garanzha, Igor Kaporin, Liudmila Kudryavtseva, François Protais, Nicolas Ray, and Dmitry Sokolov. 2021. Foldover-free maps in 50 lines of code. *ACM Transactions*

on Graphics 40, 4 (2021). https://doi.org/10.1145/3450626.3459847

Vladimir Garanzha and Liudmila Kudryavtseva. 2019. Hypoelastic Stabilization of Variational Algorithm for Construction of Moving Deforming Meshes. In *Optimization and Applications*, Yury Evtushenko, Milojica Jaćimović, Michael Khachay, Yury Kochetov, Vlasta Malkova, and Mikhail Posypkin (Eds.). Springer International Publishing, Cham, 497–511.

Vladimir Garanzha, Liudmila Kudryavtseva, and Sergei Utyuzhnikov. 2014. Variational method for untangling and optimization of spatial meshes. *J. Comput. Appl. Math.* 269 (2014), 24 – 41. https://doi.org/10.1016/j.cam.2014.03.006

Sergei Konstantinovich Godunov, Valerii Mikhailovich Gordienko, and Gennadii Aleksandrovich Chumakov. 1994. Quasi-isometric parametrization of a curvilinear quadrangle and a metric of constant curvature. *Matematicheskie Trudy* 26 (1994), 3–19.

Demetrius Gravé. 1911. Démonstration d'un théorème de Tchébychef généralisé (in French). *Journal für die reine und angewandte Mathematik* 140 (1911), 247–251.

Magnus R. Hestenes and Eduard Stiefel. 1952. Methods of conjugate gradients for solving linear systems. *Journal of research of the National Bureau of Standards* 49 (1952), 409–435.

K. Hormann and G. Greiner. 2000. MIPS: An Efficient Global Parametrization Method. In *Curve and Surface Design*. Vanderbilt University press.

Weizhang Huang. 2001. Variational Mesh Adaptation: Isotropy and Equidistribution. *J. Comput. Phys.* 174, 2 (dec 2001), 903–924. https://doi.org/10.1006/jcph.2001.6945

S.A. Ivanenko. 2000. Control of cell shapes in the course of grid generation. *Zh. Vychisl. Mat. Mat. Fiz.* 40 (01 2000), 1662–1684.

Olivier-P Jacquotte. 1988. A mechanical model for a new grid generation method in computational fluid dynamics. *Computer methods in applied mechanics and engineering* 66, 3 (1988), 323–338.

Shahar Z. Kovalsky, Noam Aigerman, Ronen Basri, and Yaron Lipman. 2014. Controlling Singular Values with Semidefinite Programming. 33, 4 (2014). https://doi.org/10.1145/2601097.2601142

Shahar Z. Kovalsky, Noam Aigerman, Ronen Basri, and Yaron Lipman. 2015. Large-scale bounded distortion mappings. *ACM Transactions on Graphics (proceedings of ACM SIGGRAPH Asia)* 34, 6 (2015).

Zohar Levi and Denis Zorin. 2014. Strict Minimizers for Geometric Optimization. *ACM Trans. Graph.* 33, 6, Article 185 (nov 2014), 14 pages. https://doi.org/10.1145/2661229.2661258

Yaron Lipman. 2012. Bounded Distortion Mapping Spaces for Triangular Meshes. *ACM Trans. Graph.* 31, 4, Article 108 (July 2012), 13 pages. https://doi.org/10.1145/2185520.2185604

John Milnor. 1969. A Problem in Cartography. *The American Mathematical Monthly* 76, 10 (1969), 1101–1112. http://www.jstor.org/stable/2317182

William Prager. 1957. On ideal locking materials. *Transactions of the Society of Rheology* 1, 1 (1957), 169–175.

Michael Rabinovich, Roi Poranne, Daniele Panozzo, and Olga Sorkine-Hornung. 2017. Scalable Locally Injective Mappings. *ACM Trans. Graph.* 36, 2, Article 16 (April 2017), 16 pages. https://doi.org/10.1145/2983621

Yu G Reshetnyak. 1966. Bounds on moduli of continuity for certain mappings. *Siberian Mathematical Journal* 7, 5 (1966), 879–886.

Martin Rumpf. 1996. A variational approach to optimal meshes. *Numer. Math.* 72, 4 (1996), 523–540.

Christian Schüller, Ladislav Kavan, Daniele Panozzo, and Olga Sorkine-Hornung. 2013. Locally Injective Mappings. *Computer Graphics Forum (proceedings of Symposium on Geometry Processing)* 32, 5 (2013).

Olga Sorkine and Marc Alexa. 2007. As-Rigid-as-Possible Surface Modeling. In *Proceedings of the Fifth Eurographics Symposium on Geometry Processing* (Barcelona, Spain) *(SGP '07)*. Eurographics Association, Goslar, DEU, 109–116.

O. Sorkine, D. Cohen-Or, R. Goldenthal, and D. Lischinski. 2002. Bounded-distortion piecewise mesh parameterization. In *IEEE Visualization, 2002. VIS 2002.* 355–362. https://doi.org/10.1109/VISUAL.2002.1183795

Jian-Ping Su, Xiao-Ming Fu, and Ligang Liu. 2019. Practical Foldover-Free Volumetric Mapping Construction. *Computer Graphics Forum* 38, 7 (2019), 287–297. https://doi.org/10.1111/cgf.13837 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13837

P.L. Tchebychev. 1853. *Théorie des mécanismes connus sous le nom de parallélogrammes.* Imprimerie de l'Académie impériale des sciences.

X Xu, Weizhang Huang, R. Russell, and J. Williams. 2011. Convergence of de Boor's algorithm for the generation of equidistributing meshes. *IMA J. Numer. Anal.* 31 (04 2011), 580–596. https://doi.org/10.1093/imanum/drp052

C. Zhu, R. H. Byrd, and J. Nocedal. 1997. L-BFGS-B: Algorithm 778: L-BFGS-B, FORTRAN routines for large scale bound constrained optimization. *ACM Trans. Math. Software* 23, 4 (1997), 550–560.