



HAL
open science

Practical lowest distortion mapping

Vladimir Garanzha, Igor Kaporin, Liudmila Kudryavtseva, François Protais,
David Desobry, Dmitry Sokolov

► **To cite this version:**

Vladimir Garanzha, Igor Kaporin, Liudmila Kudryavtseva, François Protais, David Desobry, et al..
Practical lowest distortion mapping. 2022. hal-03546749v1

HAL Id: hal-03546749

<https://hal.science/hal-03546749v1>

Preprint submitted on 28 Jan 2022 (v1), last revised 2 Nov 2022 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Practical lowest distortion mapping

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

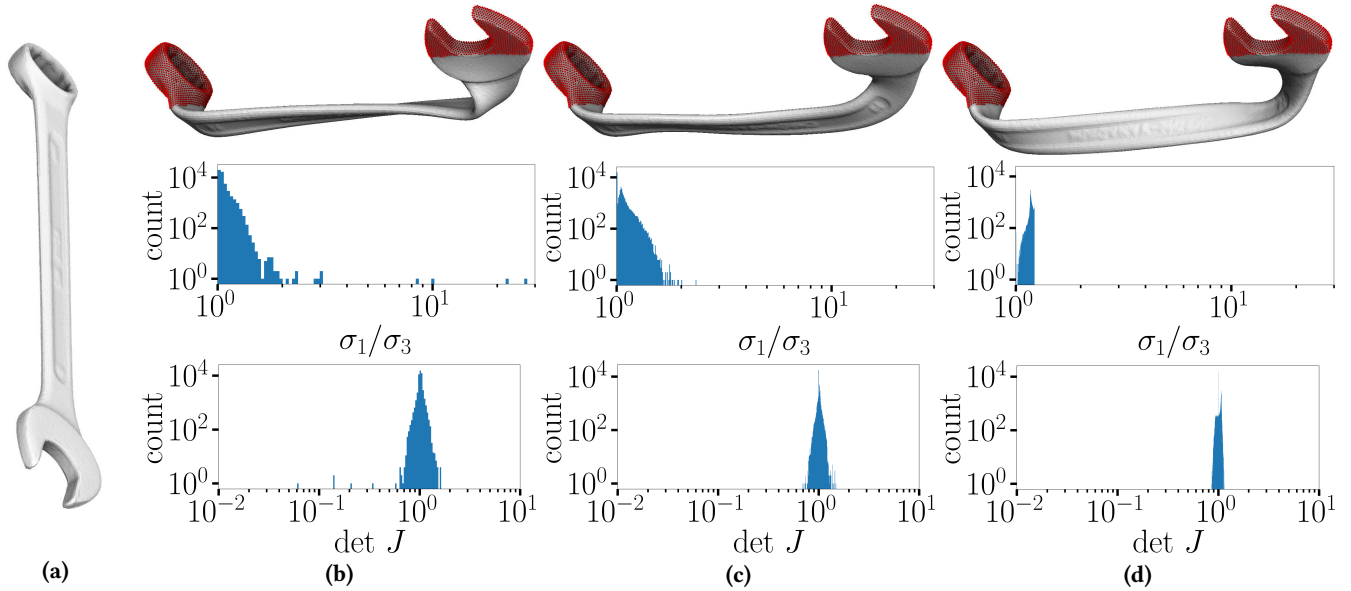


Fig. 1. Tetrahedral mesh deformation, locked vertices are shown in red. (a): Rest shape, (b): ABCD [Naitsat et al. 2020], (c): untangling [Garanzha et al. 2021a], (d) our lowest distortion mapping with $\theta = \frac{1}{2}$. **Top row**: deformations, **middle and bottom rows**: log-log histograms of element quality (condition number of the Jacobian matrix and Jacobian determinant).

Construction of optimal deformations is one of the long standing problems of computational mathematics. We consider the problem of computing quasi-isometric deformations with minimal possible quasi-isometry constant (global estimate for relative length change).

We build our technique upon [Garanzha et al. 2021a], a recently proposed numerical optimization scheme that provably untangles 2D and 3D meshes with inverted elements by partially solving a finite number of minimization problems. In this paper we show the similarity between continuation problems for mesh untangling and for attaining prescribed deformation quality threshold. Both problems can be solved by a finite number of partial solutions of optimization problems which are based on finite element approximations of parameter-dependent hyperelastic functionals. Our method is based on a polyconvex functional which admits a well-posed variational problem.

To sum up, we reliably build 2D and 3D mesh deformations with smallest known distortion estimates (quasi-isometry constants) as well as stable quasi conformal parameterizations for very stiff problems.

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

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

*Corresponding author: dmitry.sokolov@loria.fr

1 INTRODUCTION

Construction of optimal deformations is one of central themes in mesh generation research. Generally, for computing mesh deformations, using elasticity analogy was found to be very fruitful and resulted in efficient engineering mesh generation algorithms with sound theoretical foundations [Jacquotte 1988; Rumpf 1996]. The idea is to say that a mesh represents an elastic material, whose stored energy of deformation can be measured as $\int_{\Omega} f(J)$, where Ω is the input domain, J is the Jacobian matrix of elastic deformation, and f is a measure of distortion. Then, obviously, we want to minimize the stored energy of deformation.

More precisely, let us say that we want to compute a map $\vec{x}(\vec{\xi}) : \Omega \subset \mathbb{R}^d \rightarrow \mathbb{R}^d$, where d stands for the dimension (2 or 3), and the arrow denotes a d -dimensional vector. Consider the following variational problem:

$$\arg \min_{\vec{x}(\vec{\xi})} \int_{\Omega} f(J) d\xi, \quad (1)$$

where J is again the Jacobian matrix of the mapping $\vec{x}(\vec{\xi})$.

While this formulation allows to minimize distortion on average, it does not allow to limit maximum distortion. The problem of

constructing bounded distorted deformations has a long-standing history in elasticity research and goes back to 1957 Prager’s work on “Ideal locking materials” [Prager 1957]. Now this problem is referred to as “stiffening” of elastic material and generally is formulated as a set of nonlinear constraints on the Jacobian matrix of elastic deformation [Ciarlet and Necas 1985].

Constructing elastic deformations with bounded distortion constraints is a very hard non-convex and non-linear problem. There were numerous attempts made to solve the problem. 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 and Liu 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. 2015] alternate between energy optimization and a non-trivial projection to the highly non-convex set of constraints. [Lipman 2012] formulates the problem as a second-order cone programming, relying on elaborated commercial solvers such as MOSEK [Andersen and Andersen 2000].

All these papers try to incorporate the boundedness constraint into different black-box optimization toolboxes. While the approach may work reasonably well in practice, it is hard to obtain any guarantees, and solutions may exhibit undesirable, hard to explain and eliminate artifacts, say, noise and loss of symmetries. We propose to explore another research direction based on a **unconditional** optimization, avoiding altogether all issues related to constraints.

A very interesting idea [Garanzha 2000] is to consider a quasi-isometric hyperelastic material, which unlike other known models, provides admissible deformations with bounded global distortion (bounded quasi-isometry constant) as minimizers of elastic energy. Invertibility theorem for deformation of this material was established in 2D and 3D cases [Garanzha et al. 2014]. The main idea is to use controlled stiffening of material which is incorporated directly into definition of the density of deformation energy in such a way that when local measure of deformation exceeds certain threshold, the elastic material becomes infinitely stiff. The stiffening threshold is introduced as a parameter, and max-norm optimization problem for deformation is formulated as a continuation problem for polyconvex functional (minimization of stiffening threshold, or, alternatively, maximization of the quality threshold). Unfortunately, this work lacks a robust strategy of stiffening parameter choice.

Generally speaking, the stiffening technique may be expressed in the following way. Having a distortion measure $f(J)$, we can try to minimize the following energy:

$$\int_{\Omega} w \cdot f(J) d\xi,$$

where w is a weight function. Typically, w is chosen to be large in the regions where small values of $f(J)$ are required. We can use this general weighted formulation to control pointwise behaviour of the spatial distribution of the distortion measure. This idea was suggested in [Garanzha 2000] to build orthogonal mappings near boundaries which is one of the key requirements for CFD meshes. [Bommes et al. 2009] used such weights 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].

Some other methods have also used this technique. For example, by setting

$$w := f^p(J), \quad p > 0,$$

one can get a power law enhancement, thus penalizing large values of local distortion [Garanzha and Kudryavtseva 2019]. The same idea was used in IDP algorithm [Fang et al. 2021] where the authors suggested to use $p = 4$ as a rule of thumb. If function f is polyconvex, function $f^{1+p}(J)$ is polyconvex as well. Note however, that in the continuous case one cannot prove that $f^{1+p}(J)$ is bounded from above, and in the discrete case one cannot prove that $f(J)$ does not grow to infinity under mesh refinement.

Another weight

$$w := \frac{1}{1 - tf(J)}, \quad 0 \leq t < 1$$

corresponds to algorithm from [Garanzha 2000]. The resulting functional has an infinite barrier (refer to §2.3 for more details) on the boundary of the set of quasi-isometric deformations [Garanzha et al. 2014], thus solving problem of quasi-isometric map generation formulated by Godunov [S. K. Godunov 1995] for general domains.

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. (16)). With this new contribution, we were able to confirm the 20-years old conjecture that variational problem [Garanzha 2000] allows to build best deformations compared to state-of-the-art algorithms.

Our contributions. We propose a very simple algorithm that allows us to reliably build 2D and 3D mesh deformations with **smallest known distortion estimates** (quasi-isometry constants). To the best of our knowledge, we are the first to provide **theoretical guarantees**

to this long standing problem. Our 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. We show that attainable quality threshold estimates (quasi-isometry constants) **do not deteriorate under mesh refinement** which is a unique property of the proposed algorithm.

By coupling our technique with [Garanzha et al. 2021a], we obtain a complete unified mapping pipeline. For a better reproducibility, we publish a complete **C++ implementation** [Sokolov 2022]. We start from an arbitrary initial deformation, untangle the mesh in a finite number of steps, minimizing mean distortion, and finally we minimize the maximum distortion. Just like for the untangling step, we can obtain a deformation with a prescribed quality threshold in a finite number of steps. 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.

Last, but not least, we bring more robustness to global parameterizations. [Garanzha et al. 2021a] produce maps free of inverted

elements, but do not prevent k -coverings, thus losing local injectivity. We provide a very simple but effective way to handle this problem: we guarantee **local injectivity for global parameterizations**.

2 VARIATIONAL FORMULATION FOR UNTANGLING AND DISTORTION MINIMIZATION

This section gives a necessary background on elastic deformations. First, in §2.1 we revisit main issues of mesh deformation based on the elasticity theory. Then, in §2.2 we present the core idea behind the untangling procedure described in [Garanzha et al. 2021a]. Next, in §2.3, we describe the idea behind generation of deformations of a given quality, until recently very heuristic. Finally, having prepared all necessary concepts, we can present our latest guarantees and results (§3 and §4).

2.1 Elastic material choice and main issues

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; Hormann and Greiner 2000]:

$$f(J) := (1 - \theta)f_s(J) + \theta f_v(J), \quad (2)$$

where shape distortion is defined as

$$f_s(J) := \begin{cases} \frac{1}{d} \frac{\text{tr} J^T J}{(\det J)^{\frac{2}{d}}}, & \det J > 0 \\ +\infty, & \det J \leq 0 \end{cases} \quad (3)$$

while volumetric distortion is defined

$$f_v(J) := \begin{cases} \frac{1}{2} \left(\det J + \frac{1}{\det J} \right), & \det J > 0 \\ +\infty, & \det J \leq 0 \end{cases} \quad (4)$$

Note that functions $f_s(J)$ and $f_v(J)$ have concurrent goals, one preserves angles and the other preserves volumes, and thus θ serves as a trade-off parameter. Density (2) is a polyconvex function satisfying ellipticity conditions, it is therefore very well suited for a numerical optimization.

Polyconvexity of the energy, ellipticity of the PDE and invertibility of the minimizer are very strong results; moreover, the minimizer of Prob. (1) has a minimum average distortion. There are, however, two issues with this formulation:

(a) the variational problem makes sense and allows for minimization only when an **initial guess** is in the **admissible** domain $\min_{\Omega} \det J > 0$, so a special untangling procedure is required for an arbitrary initial guess;

(b) the fact that functional (1) is bounded **does not imply** that distortion measure (2) is **bounded**. It optimizes average deformation and admits integrable singularities. In order to provably suppress this behaviour, one has to consider modified hyperelastic material which forbids deformations with local distortion above prescribed threshold.

Two following sections discuss both points and lay the ground for our contribution in controlled stiffening of a hyperelastic material, that provably allows us to build best known quasi-isometric maps.




Fig. 2. Regularization function for the denominator in Eq. (6). When ε tends to zero, $\chi(\varepsilon, D)$ tends to $\max(0, D)$.

2.2 Untangling

With a slight abuse of notations, the density (2) can be rewritten as follows:

$$f := (1 - \theta) \frac{1}{d} \frac{\text{tr} J^T J}{(\max(0, \det J))^{\frac{2}{d}}} + \theta \frac{1}{2} \frac{1 + \det^2 J}{\max(0, \det J)}$$

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 function $\chi(D, \varepsilon)$ that will serve as a regularization for $\max(0, D)$ in the denominator:

$$\chi(D, \varepsilon) := \frac{D + \sqrt{\varepsilon^2 + D^2}}{2} \quad (5)$$

When ε tends to zero, $\chi(\varepsilon, D)$ tends to $\max(0, D)$, refer to Fig. 2 for an illustration. Then, the density can be regularized as follows [Garanzha and Kaporin 1999]:

$$f_\varepsilon(J) := (1 - \theta) \frac{1}{d} \frac{\text{tr} J^T J}{(\chi(\det J, \varepsilon))^{\frac{2}{d}}} + \theta \frac{1}{2} \frac{1 + \det^2 J}{\chi(\det J, \varepsilon)} \quad (6)$$

Finally, Prob. (1) can be rewritten as follows:

$$\lim_{\varepsilon \rightarrow 0^+} \arg \min_{\vec{x}(\xi)} \int_{\Omega} f_\varepsilon(J) d\xi \quad (7)$$

This formulation suggests an algorithm: build a decreasing sequence of the $\varepsilon^k \rightarrow 0$, and for each value ε^k solve an optimization problem. In other words, Prob. (7) 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. (7) 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 $X := (\vec{x}_1^T \dots \vec{x}_{\#V}^T)^T$.

A discretization of Prob. (7) can be written as follows:

$$\lim_{\varepsilon \rightarrow 0^+} \arg \min_X F(X, \varepsilon), \quad (8)$$

$$\text{where } F(X, \varepsilon) := \sum_{k=1}^{\#T} f_\varepsilon(J_k) \text{vol}(T_k),$$

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

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** $\min_{t \in \{1, \dots, \#T\}} \det J_t^k > 0$ 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

Prob. (8) can be solved with Alg. 1. The algorithm itself is very simple, and has been published more than 20 years ago [Garanzha and Kaporin 1999]. Note, however, that the crucial step here is the choice of the regularization sequence ε^k (Alg. 1–line 3), and until very recently only heuristics were used. Last year [Garanzha et al. 2021a] have presented a way to build such a sequence that offers theoretical guarantees on untangling (refer to §3 for a complete formulation).

Function $F(X, 0)$ has an impenetrable infinite barrier on the boundary of the set of meshes with positive cell volumes¹

$$\frac{\text{vol}(\vec{x}(T_k))}{\text{vol}(T_k)} > 0, \quad k = 1, \dots, \#T \quad (9)$$

which is a finite-dimensional approximation of the set $\det J > 0$. Untangling in Prob. (8) is guaranteed because [Garanzha et al. 2021a] build a decreasing sequence $\varepsilon^k \rightarrow 0$ that forces the mesh to fall into the feasible set (9) of untangled meshes. With some assumptions on the minimization toolbox chosen, untangling is guaranteed to succeed in a finite number of steps.

2.3 Controlled quasi-isometric stiffening (QIS): minimization of maximum distortion

In addition to untangling, by solving Prob. (8) we minimize **average** distortion of a map. In this section we present the idea behind [Garanzha 2000] that will allow us to compute a deformation with prescribed quality, i.e. optimize the **maximum** distortion until it reaches a given bound.

Consider following variational problem related to construction of deformations with prescribed quality $0 \leq t^* < 1$:

$$\arg \min_{\vec{x}(\vec{\xi})} \int_{\Omega} \frac{f(J)}{1 - t^* f(J)} d\xi, \quad (10)$$

¹As a side note, this set has a quite complicated structure. For k -th simplex $\text{vol}(\vec{x}(T_k))$ is a polylinear function of coordinates of its vertices, hence each term in (9) defines a non-convex set. One can hardly expect that intersection of the sets in (9) would result in a convex domain. Moreover, Ciarlet [Ciarlet and Geymonat 1982] has proved that 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 // 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. (16)
- 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)

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

$$f(J) < \frac{1}{t^*} \quad (11)$$

Here parameter t^* plays the role of the lower quality bound of the deformation. Note that the density in Prob. (10) is polyconvex and this variational problem is well-posed [Garanzha et al. 2014].

We propose following finite-dimensional approximation of Prob. (10):

$$\lim_{t \rightarrow t^*} \arg \min_X W(X, t), \quad (12)$$

$$\text{where } W(X, t) := \sum_{k=1}^{\#T} \frac{f(J_k)}{1 - t f(J_k)} \text{vol}(T_k),$$

It is important to note that a solution of Prob. (8) corresponds to a solution of Prob. (12) with $t^* = 0$, i.e. when no bound on the maximum deformation is imposed. But then, having reached the set (9), we can build an increasing sequence of $t^k \rightarrow t^*$ to contract the set until the mesh falls into the set of deformations with prescribed quality t^* :

$$f(J_k) < \frac{1}{t^*}, \quad k = 1, \dots, \#T \quad (13)$$

Alg. 2 sums up 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.

While we assume that parameter t^* exists, fortunately we are not obliged to know it to make QIS algorithm work. It is an important advantage over optimization algorithms which use prescribed distortion bound like LBD [Kovalsky et al. 2015], since in practice even rough estimates of this bound are not available. Essentially, QIS algorithm by itself serves as a distortion bound estimation tool for problems of any complexity. Refer to App. A for the argumentation of the fact that minimizing our distortion measure implies minimization of the upper bound for the quasi-isometry constant.

3 THEORETICAL GUARANTEES FOR QUASI-ISOMETRIC STIFFENING

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. 2021a, 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 (9) 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

$$F(X^{k+1}, \varepsilon^k) \leq (1 - \sigma)F(X^k, \varepsilon^k), \quad (14)$$

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

$$\min_X F(X, \varepsilon^k) > (1 - \sigma)F(X^k, \varepsilon^k). \quad (15)$$

Then the feasible set (9) 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. 2021a, 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} :

$$t^{k+1} := t^k + \sigma \frac{1 - t^k f_+^{k+1}}{f_+^{k+1}}, \quad (16)$$

where $0 < \sigma < 1$ is again the performance index of the minimization toolbox. Clearly, formula (16) 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 (13) 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

$$W(X^{k+1}, t^k) \leq (1 - \sigma)W(X^k, t^k), \quad (17)$$

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

$$\min_X W(X, t^k) > (1 - \sigma)W(X^k, t^k). \quad (18)$$

Then the feasible set (13) 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 $W(X^K, t^K) < +\infty$.

PROOF. The main idea is very simple: update rule (16) 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 (13) 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, what contradicts the boundedness.

To prove that $\{W(X^k, t^k)\}_{k=0}^\infty$ is bounded from above, we analyse the behavior at some iteration k . First of all, if we couple update rule (16) with the fact that for any $t_2 > t_1$ the ratio $\frac{1-t_1\psi}{1-t_2\psi}$ is increasing function of argument ψ , we can see that the following inequality holds:

$$(1 - \sigma)W(X^{k+1}, t^{k+1}) \leq W(X^{k+1}, t^k). \quad (19)$$

More precisely,

$$\begin{aligned} W(X^{k+1}, t^{k+1}) &= \sum_i \frac{1 - t^k f_i(X^{k+1})}{1 - t^{k+1} f_i(X^{k+1})} \frac{f_i(X^{k+1})}{1 - t^k f_i(X^{k+1})} \leq \\ &\sum_i \frac{1 - t^k f_+(X^{k+1})}{1 - t^{k+1} f_+(X^{k+1})} \frac{f_i(X^{k+1})}{1 - t^k f_i(X^{k+1})} = \frac{1}{1 - \sigma} W(X^{k+1}, t^k). \end{aligned}$$

Then, at each iteration k , either condition (17) or condition (18) must be satisfied. Let us consider both cases.

Cond. (17) holds: in this case function W actually decreases. Eq. (19) combined with Cond. (17) directly imply that

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

Cond. (18) holds: under assumption that $t^{k+1} < t^*$, by combining Eq. (19) and Cond. (18), we obtain:

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

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 (16) 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

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

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

REMARK 1. An important corollary of Th. 2 is that, provided that the admissible set (13) 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 Th. 2, therefore it can reach the distortion bound in a finite number of steps.

In practice, just like in [Garanzha et al. 2021a], the global estimate σ is not known in advance. Garanzha et al. suggest to compute the local descent coefficient for each minimization step, and so do we. Instead of σ in the update rule (16), 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).

4 RESULTS AND DISCUSSION

Recall that our contribution is two-fold: **(a)** we provide an algorithm to compute a deformation where all cells have a deformation quality above given threshold, and **(b)** we guarantee local invertibility of a global parameterization for quad generation. This section is organized in two subsections accordingly.

4.1 Quality optimization results

In order to check the ability of the variational method (12) to improve worst quality elements in the deformed meshes, we have performed a series of tests.

Constrained boundary injective mapping. Along with their paper [Du et al. 2020], Du et al. have published a benchmark with a huge number of 2D and 3D constrained boundary injective mapping challenges. We have computed initial injective maps with [Garanzha et al. 2021a], and optimized the quality of the maps. Fig. 3 shows the scatter plot for every mapping challenge from the database. For each problem the plot has two dots: the red one corresponds to the input injective map reduced to two numbers, namely, the maximum stretch and the minimum scale. The green dot is the quality of the map after our optimization. Gray segments show the correspondence between the dots. As can be seen from the plot, in the vast majority of cases our optimization improves both the maximum stretch and minimum scale quality measures despite the fact that the boundary is locked.

Conformal mapping. While in theory conformal maps are beyond the scope of QIS algorithm, in practice it was found to be highly successful. As was suggested in [Garanzha et al. 2021a], approximation of conformal mappings with quasi-isometric ones is indeed a good and stable numerical solution. We have computed six discrete conformal maps for a triangular mesh of a hand (refer to Fig.4). To compare quality of the maps, we use the condition number of the Jacobian matrix $\frac{\sigma_1(J)}{\sigma_2(J)}$, where σ_1 and σ_2 stand for the (ordered) singular values of J . First we have computed four maps with state-of-the-art methods:

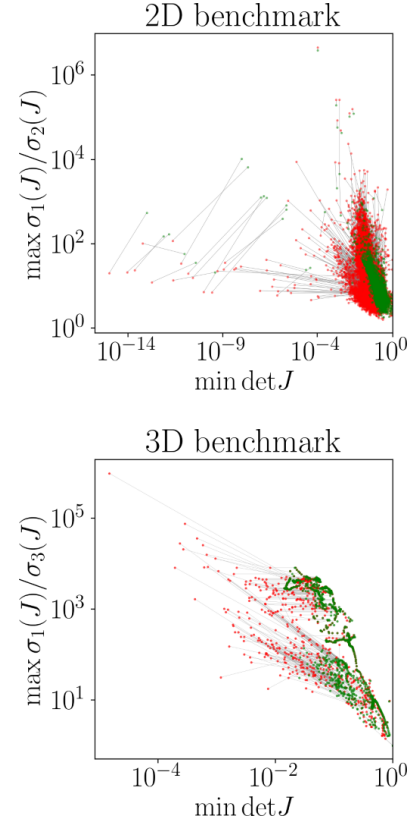


Fig. 3. Quality plot of the resulting locally injective maps for every challenge from the database provided by [Du et al. 2020]. Each dot corresponds to a quality of the corresponding map reduced to two numbers: the maximum stretch and the minimum scale. Untangling results [Garanzha et al. 2021a] are shown in red, whereas our results are shown in green (we took $\theta = \frac{1}{2}$). The gray lines show the correspondence in the results. **Top**: mapping quality on the 2D dataset (10743 challenges). **Bottom**: mapping quality on the 3D dataset (904 challenges).

(a) The easiest one to compute is the least squares conformal map [Lévy et al. 2002]. This well-known method requires solving one linear system with a symmetric positive definite matrix. The idea behind LSCM is to compute a P^1 finite element approximation of the Cauchy-Riemann conditions over all triangles of the mesh.

(b) Second map was obtained by applying the boundary first flattening method [Sawhney and Crane 2017] with boundary log-scale factors set to zero. This choice of boundary condition leads to the conformal map with minimal area distortion [Springborn et al. 2008, App. E].

(c) Third map is the result of the elliptic smoother [Garanzha et al. 2021a], we have obtained it by solving Eq. (10) with $\theta = 0$, $t = 0$.

(d) Fourth map was obtained with CEPS [Gillespie et al. 2021]. Note that this method can alter the input triangulation, and the way it handles surfaces with boundary is to glue together two copies of the input mesh along the boundary, introduce cone singularities and cut the mesh. The seams are shown in red in Fig. 4-(e).




Fig. 4. Discrete conformal mapping test. **(a)**: LSCM [Lévy et al. 2002], **(b)**: BFF [Sawhney and Crane 2017], **(c)**: untangling [Garanzha et al. 2021a], **(d)**: CEPS [Gillespie et al. 2021], **(e)**: QIS applied to the untangling result (see **(c)**), **(f)**: QIS applied to the surface cut by CEPS (see **(e)**).

Finally, starting from the untangling **(c)** and CEPS **(d)**, we have computed optimal discrete conformal maps by solving Eq. (10) with $\theta = 0$ while maximizing for t . The results are shown in Fig. 4-**(e)** and Fig. 4-**(f)**. Log-log element quality histograms show that we improve considerably the quality of input maps. It is easy to see that the maximum condition number of the Jacobian matrix is consistently better in our maps (1.33 and 1.28, respectively).

We were quite surprised by the fact that QIS algorithm outperforms highly elaborated algorithms for conformal parameterizations, since this problem is beyond its derivation principles. One possible explanation, which is referred to in [Garanzha et al. 2021a], is the

hypothesis that approximation of conformal mappings by quasi-isometric mappings with growing quasi-isometry constants is a good theoretical and engineering way for conformal parameterizations.

Comparison with ABCD. Fig. 1 shows a comparison with Adaptive Block Coordinate Descent for Distortion Optimization [Naitsat et al. 2020]. We took a tetrahedral mesh of a combination wrench, and we imposed positional constraints on the vertices located on both ends of the wrench. As before, we optimize the quality starting from the untangling result (Fig. 1-**(c)**).

While the measures of deformation differ (ABCD uses ARAP-like energy), most of our input map elements have comparable quality to ABCD. Note however that ABCD produces several elements of very bad quality, and untangling gains an order of magnitude over the worst element quality: the maximum stretch $\max \sigma_1/\sigma_3 \approx 28$, the minimum scale $\min \det J \approx 0.06$ for ABCD and $\max \sigma_1/\sigma_3 \approx 2.35$, $\min \det J \approx 0.69$ for the untangling. In our turn, we improve the quality even further, our quality measures are $\max \sigma_1/\sigma_3 \approx 1.22$ and $\min \det J \approx 0.86$.

Comparison with Injective Deformation Processing. Fang et.al [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 exponential factor $p > 1$.

In our next test we deform a cylinder tetrahedral mesh. We applied two bone handles (two thin boxes of interior axis vertices) to bend it. Fig. 5 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 w.r.t regular IDP, but does not allow to eliminate it completely. As can be seen in Fig. 5-left, the stress is concentrated around the locked vertices (shown in magenta). Our optimization ($\theta = \frac{1}{2}$) allows 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.

Comparison with 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.

In our test (refer to Fig. 6), 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 a particular interest. In our test the minimum Jacobian determinant increases from 0.30 to 0.69, while the maximum stretch *increases* from 1.50 to 1.61. Despite the increase in the max stretch, our variational problem leads to a much better overall element quality distribution. Note that this test was performed with default stretch/scale trade-off parameter $\theta = \frac{1}{2}$. If we choose, for example, $\theta = \frac{1}{3}$, both quality measures improve: $\max \sigma_1/\sigma_2 \approx 1.43$ and $\min \det J \approx 0.76$.

Comparison with Simplex Assembly. Our final test is confronting our optimization 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. Having disassembled the mesh, the simplices are then assembled by minimizing the mapping distortion, while keeping the mapping




Fig. 5. 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 log-log histograms and corresponding color plots.




Fig. 6. Comparison of LBD [Kovalsky et al. 2015] vs QIS. **(a):** 3D surface to flatten is a regular triangular mesh of a square patch that was lifted and noised. **(b):** The map obtained by LBD. **(c):** QIS ($\theta = \frac{1}{2}$). **Top row:** flattenings of **(a)**, colors correspond to the quality of elements (conditioning of the Jacobian). **Middle and bottom rows:** log-log element quality histograms.




Fig. 7. Two quasi-isometric maps for the “Lucy” mesh. (a): Simplex assembly, (b): QIS, $\theta = \frac{1}{2}$. The quality of elements and its distribution over the surface is shown in log-log histograms and corresponding color plots.




Fig. 8. Mapping a half-sphere on a disc: our solution with $\theta = \frac{1}{2}$ (a) vs Simplex Assembly (b). **Top row:** condition number of the Jacobian matrix, **bottom row:** Jacobian determinant.

feasible. Fig. 7 provides a quality comparison of SA with our quasi-isometric ($\theta = \frac{1}{2}$) map for a free-boundary mapping of the “Lucy” mesh. This comparison is interesting for two different reasons: first, SA offers an explicit optimization for the distortion bound, and second, in 2D Fu et. al use exactly the same distortion measure as we do.

Our method shows consistently better maps w.r.t. Simplex Assembly. In this test, the worst condition number $\max \frac{\sigma_1(J)}{\sigma_2(J)}$ is 16.46 for SA and 10.32 for our method. The minimum scale $\min \det J$ is equal to 0.05 for SA and 0.10 for our method.

Note that while in 2D Simplex Assembly has the same objective function as our method, the way SA poses the problem (minimization of maximal distortion) leads to non-smooth solutions. The fact that the solution is noisy can already be seen in Fig. 7, but we chose to perform a second test to highlight the fact: we map a half-sphere onto a disc (Fig. 8).

This test case is interesting because it has a closed form solution. It is possible to build an analytical flattening which has the smallest known quasi-isometry constant $\Gamma = \sqrt{\pi/2}$. This mapping can be obtained by isometric projection of meridians onto straight segments on the plane starting from the north pole while keeping angular projection uniform. Obviously, singular values of this mapping range from 1 to $\pi/2$, and using best scaling we get $\Gamma = \sqrt{\pi/2} \approx 1.253$.

As before, our result is better: the distortion for our flattening is equal to $\sqrt{\frac{\max_i \sigma_1(J_i)}{\min_i \sigma_2(J_i)}} = 1.279$, which is within 2% from the ideal bound, and the quasi-isometry constant for SA is equal to 1.47. This difference is due to the optimization scheme choice. Since we discretize a well-posed variational problem, our method provides smooth solutions, whereas SA result is noisy and loses angular symmetry.

Stability. As we have already mentioned, our 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. Here we show that attainable quality threshold estimates (quasi-isometry constants) do not deteriorate with mesh coarsening which is unique

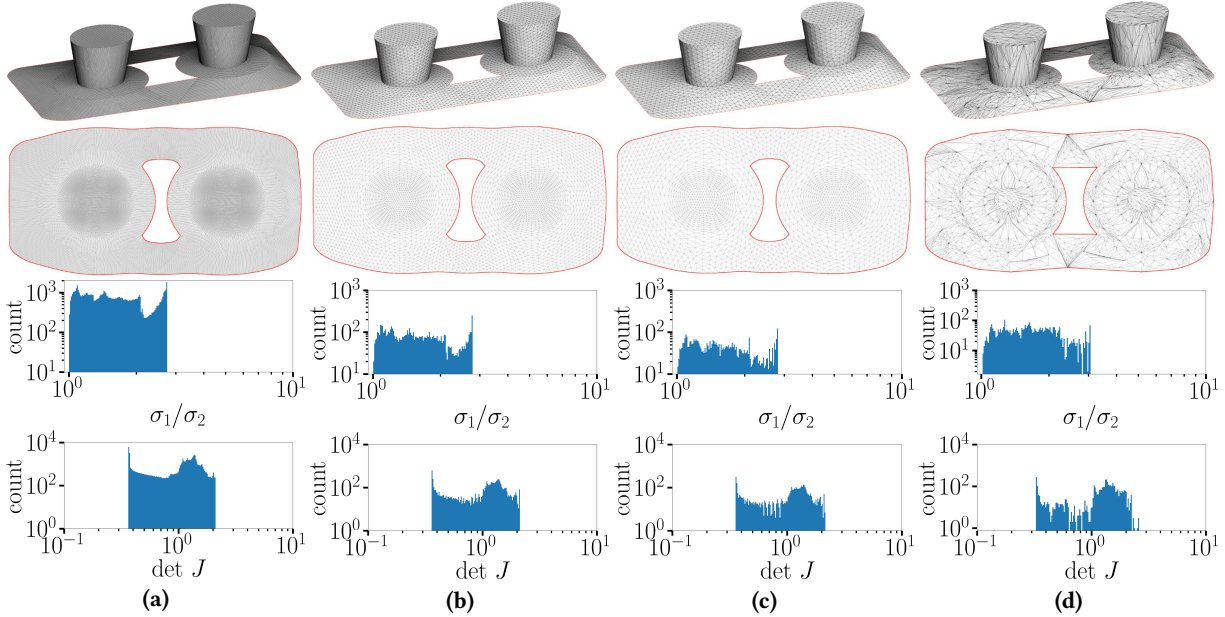


Fig. 9. Our method shows an excellent stability irrespective of input mesh quality. Here we show four free-boundary quasi-isometric maps ($\theta = \frac{1}{2}$) computed on different meshes. (a)–(c): Behavior under coarsening (from 73k to 4k triangles) of an isotropic mesh. (d): Stress test, the input mesh has a very bad quality with the maximum aspect ratio of 10^8 over the triangulation. Despite that, the resulting map is still of a good quality.

property of the proposed algorithm. To illustrate this point, we have computed four free-boundary quasi-isometric maps on different meshes of the same object (Fig. 9). Under coarsening of an input isotropic mesh, the distortion bound remains almost the same (2.07, 2.10, 2.11 for the maximum jacobian condition number, respectively). We have also performed a stress test Fig. 9–(d) with maximum aspect ratio of 10^8 over the input triangulation, and the maximum jacobian condition number we obtained equals to 2.25.

4.2 Untangling global parameterizations

An important application of parameterizations is the generation of quad meshes. Given a 3D surface and its 2D flattening, applying the inverse of the map to a 2D grid generates a grid on the surface, i.e. a quad mesh. Naturally, this technique requires the map to be locally invertible, hence the importance of the untangling approach. For producing more complex quad meshes [Bommes et al. 2009], it is also possible to introduce discontinuities in the map. To generate a valid quad mesh, however, we need to impose constraints along these discontinuities: the transition function that maps one side of a cut to the other side must be grid-preserving, i.e. it transforms the 2D unit grid onto itself (see Fig. 10).

All such grid preserving transition functions can be decomposed into a rotation of $k\pi/2$ plus an integer translation. In practice, producing such a global parameterization requires two parameterization steps: one where only the rotation is known (from a frame field e.g. [Desobry et al. 2021]) and one where the translation is also known (after quantization e.g. [Bommes et al. 2013]). In both cases, these boundary constraints are affine and can be easily introduced




Fig. 10. The map is discontinuous across the red cut, but the projection of the unit grid cells from the map (right) to the object (left) coincide thanks to the grid preserving transition function (green arrow).

in our optimization scheme using [Bommes et al. 2012]. Figure 11 illustrates the global parameterization pipeline.

Recall that being free of inverted elements does not imply local invertibility [Garanzha et al. 2021a, §3.3]. In difficult cases, double coverings may appear. It is possible to prevent this with a brute force solution [Garanzha et al. 2021b]. Local injectivity can be enforced by adding extra (virtual) triangles. Unfortunately, this approach rigidifies the mesh and can be time consuming. We can do better for global parameterizations.

We can interpret the gradients of the parameterization as a frame field [Ray et al. 2008]. Double coverings arise from index -1 singularities of this field (refer to Fig. 12). Poincaré-Hopf theorem states that the sum of indices is constant, so an index 1 singularity must be placed somewhere to compensate for the index -1. With free boundary mapping (as in our example in Fig. 12), the index 1 is placed outside of the domain, simply adding a loop to the boundary. For the global parameterization case, there is no free boundary, so the index 1 must appear at a vertex. Recall however, that our maps

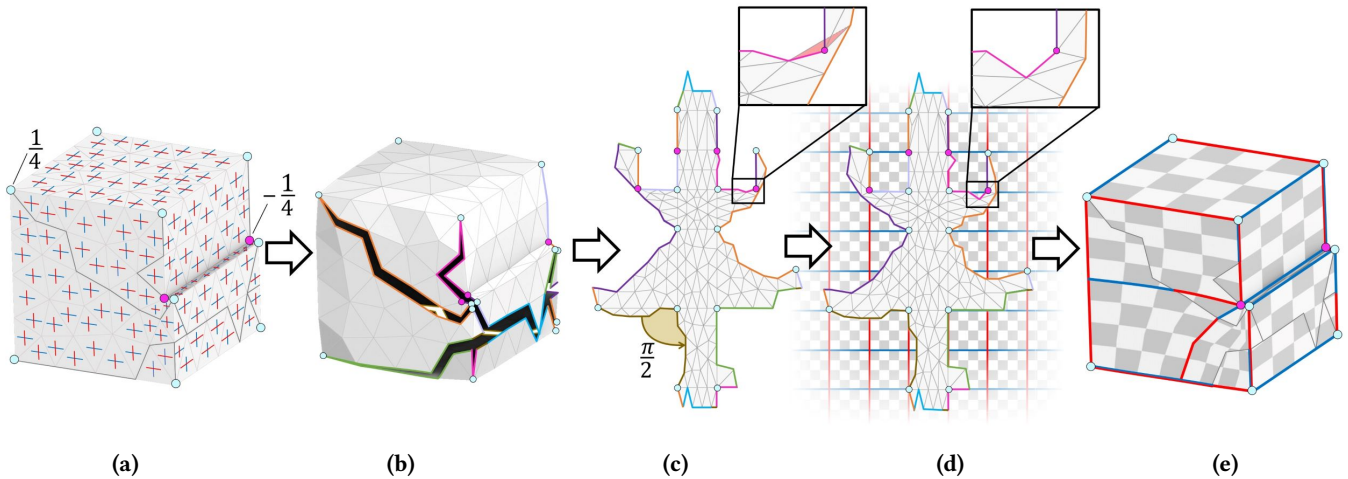


Fig. 11. Quad generation via global parameterization pipeline. **(a)**: we compute a frame field over the input triangulation, it allows us to determine singular vertices. Then the mesh is cut open **(b)** and flattened under grid preserving constraints along the cuts. Least squares solution **(c)** has inverted elements that we need to untangle **(d)**. Finally, the unit grid is projected back to the mesh to define the quad mesh **(e)**.




Fig. 12. A double covering example. **Left**: a surface to flatten is made of 12 equilateral triangles. **Middle**: the surface can be mapped to a plane without inverted elements, producing a double covering. The map is not invertible in one point. **Right**: gradients of the parameterization form two orthogonal vector fields (shown in red and blue), both of them have a -1 singularity at the center.

have positive Jacobian determinant over all elements, and thus it is impossible to place singularity 1 at a regular vertex, since it is a pole singularity that would force the map to have degenerate elements. Therefore, the only possibility for the solver is to place it on a vertex that already has a negative index singularity (typically $-1/4$) as illustrated in Fig. 13.

This observation allows to avoid double coverings by simply forcing vertices with negative index singularity to preserve the index. To this end, for each such vertex v , we flatten its one ring, and compute the rotation + scale (with respect to v) that send each adjacent vertex to the next one. The rotation is then scaled according to the index of the vertex ($\times 5/4$ for the index $-1/4$), and these affine equations tying in adjacent vertices are introduced as constraints to our system. This solution forces the angle distortion to be perfectly distributed on the one ring of these vertices. A local mesh refinement is applied to prevent these new constraints to conflict.

As illustrated on Fig. 14, our method provides injective maps, even with the high distortion required to produce coarse quad meshes.




Fig. 13. Quad mesh generation via global parameterization approach. The triangular domain is mapped to a plane, the green line corresponds to a (grid preserving) discontinuity in the map. Red and blue lines correspond to the flat unit grid under the action of the inverse map. **Left**: one singularity of index $-1/4$ is present in the domain, a valid quad mesh is generated. **Right**: a double covering (-1 singularity) can appear if and only if the $-1/4$ singularity is “promoted” to $3/4$, thus total sum is still equal to $-1/4$. The map is no longer invertible (even if all elements have positive Jacobian!), leading to problems in quad mesh generation.

5 LIMITATIONS AND FUTURE WORK

Note that the “finite number of steps” theorem cannot be applied to the case when the diameter of the feasible set is zero. This may happen, for example, when the parameter t^* is set to the maximum attainable value. To handle marginal cases of small size of the feasible set, we can assign to σ_0 parameter quite small value, say 10^{-3} . For a moment it is not clear whether we can attain deformation with best possible distortion numerically.




Fig. 14. To avoid double coverings, we constrain singularities with negative index and untangle the map (left). This allows us to obtain valid parameterization even for coarsest quad layouts (right).

QIS algorithm is based on the idea of incremental contraction of the set of admissible mesh deformations, i.e. the deformations which provide finite value of QIS functional. While rigorous analysis is missing, one can hope that with contraction the diameter of the admissible set tends to zero while its boundary, which is evidently infinitely smooth, is rounded and resembles the multidimensional sphere. Evidently one should stop at some finite diameter, expecting that minimum in the contracted set converges to limiting point faster than the diameter covers to zero. Of course, much worse alternative is possible when measure of the feasible set tends to zero while its diameter is frozen near certain fixed value. In order to get certain one should obtain estimates for positive definite part of the Hessian of QIS functional, similar to that for untangling functional in [Garanzha et al. 2021a], and consider limiting values, i.e. relation between diameter of the admissible set and curvature of its boundary. This analysis is beyond the scope of present paper.

6 CONCLUSION

We formulate a set of variational problems potentially covering the complete technological chain for construction of optimal mappings and deformations with fixed as well as free boundaries. We start with the continuation problem with respect to parameter ε , this minimization allows us to compute optimal in the average deformations. Then we formulate a continuation problem for worst quality measure maximization (quasi-isometric stiffening, QIS), which retains polyconvexity and smoothness of deformation. At all stages we take care to demonstrate that finite number of basic optimization steps is enough in order to solve the problem. We illustrate performance of our algorithm with challenging 2D and 3D numerical tests. Importance of polyconvexity is underlined since some competing algorithms for mesh optimization which potentially may attain good quality criteria for deformations tend to lose deformation smoothness and symmetries even in the simplest test cases.

A RELATIONSHIP BETWEEN DISTORTION MEASURES AND THE QUASI-ISOMETRY CONSTANT

Given a deformation of a mesh, let us denote by Γ its quasi-isometry constant (maximal relative length distortion of the map). In our optimization procedure we do not optimize for Γ directly, we minimize a mesh distortion measure f instead. Let us show the relation between the two. More precisely, we want to estimate Γ through

available values of $f(J_k)$, where J_k is the Jacobian matrix of k -th 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 :

$$\frac{1}{\Gamma} \leq \sigma_j(J_k) \leq \Gamma. \quad (20)$$

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

$$f(J_k) < \frac{1}{t}, \quad t < 1 \quad (21)$$

implies (20). Naturally, for QIS scheme to work with a certain distortion measure f , we need the bound for Γ to tend to 1 as $t \rightarrow 1$. In this section we analyze two possible choices of distortion measures, namely, Eq. (2) and Symmetric Dirichlet [Schüller et al. 2013].

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

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

and

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

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

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

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

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

Indeed, Eq. (2) forces the quasi-isometry constant Γ to tend to unity when $t \rightarrow 1$, so our QIS technique is correct.

Simple estimates for Γ can be also derived for Symmetric Dirichlet (SD) distortion measure. This measure can be written as follows:

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

In this case

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

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).

REFERENCES

- 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
- 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, Marcel Campen, Hans-Christian Ebke, Pierre Alliez, and Leif Kobbelt. 2013. Integer-Grid Maps for Reliable Quad Meshing. *ACM Trans. Graph.* 32, 4, Article 98 (jul 2013), 12 pages. <https://doi.org/10.1145/2461912.2462014>
- 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>
- David Bommes, Henrik Zimmer, and Leif Kobbelt. 2012. Practical mixed-integer optimization for geometry processing. In *Proceedings of the 7th international conference on Curves and Surfaces* (Avignon, France). Springer-Verlag, Berlin, Heidelberg, 193–206.
- P.G. Ciarlet and G. Geymonat. 1982. Sur les lois de comportement en élasticité non-linéaire 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>
- David Desobry, François Protais, Nicolas Ray, Etienne Corman, and Dmitry Sokolov. 2021. Frame Fields for CAD Models. In *Advances in Visual Computing*, George Bebis, Vassilis Athitsos, Tong Yan, Manfred Lau, Frederick Li, Conglei Shi, Xiaoru Yuan, Christos Mousas, and Gerd Bruder (Eds.). Springer International Publishing, Cham, 421–434.
- 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>
- 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>
- 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. 2021a. Foldover-free maps in 50 lines of code. *ACM Transactions on Graphics* 40, 4 (2021). <https://doi.org/10.1145/3450626.3459847>
- Vladimir Garanzha, Igor Kaporin, Liudmila Kudryavtseva, François Protais, Nicolas Ray, and Dmitry Sokolov. 2021b. On Local Invertibility and Quality of Free-boundary Deformations. In *IMR 2021 - 29th International Meshing Roundtable*. Virtual, United States. <https://doi.org/10.5281/zenodo.5559040>
- Vladimir Garanzha and Liudmila Kudryavtseva. 2019. Hypoelastic Stabilization of Variational Algorithm for Construction of Moving Deforming Meshes. In *Optimization and Applications*, Yury Evtushenko, Milojeva 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>
- Mark Gillespie, Boris Springborn, and Keenan Crane. 2021. Discrete Conformal Equivalence of Polyhedral Surfaces. *ACM Trans. Graph.* 40, 4 (2021).
- 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.
- 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. 2015. Large-scale bounded distortion mappings. *ACM Transactions on Graphics (proceedings of ACM SIGGRAPH Asia)* 34, 6 (2015).
- 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>
- Bruno Lévy, Sylvain Petitjean, Nicolas Ray, and Jérôme Maillot. 2002. Least Squares Conformal Maps for Automatic Texture Atlas Generation. In *ACM SIGGRAPH conference proceedings*, ACM (Ed.).
- Alexander Naitsat, Yufeng Zhu, and Yehoshua Y Zeevi. 2020. Adaptive Block Coordinate Descent for Distortion Optimization. In *Computer Graphics Forum*, Vol. 39. Wiley Online Library, 360–376.
- William Prager. 1957. On ideal locking materials. *Transactions of the Society of Rheology* 1, 1 (1957), 169–175.
- Nicolas Ray, Bruno Vallet, Wan Chiu Li, and Bruno Lévy. 2008. N-Symmetry Direction Field Design. *ACM Trans. Graph.* 27, 2, Article 10 (may 2008), 13 pages. <https://doi.org/10.1145/1356682.1356683>
- 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.
- G. A. Chumakov S. K. Godunov, V. M. Gordienko. 1995. Variational principle for 2-D regular quasi-isometric grid generation. *International Journal of Computational Fluid Dynamics* 5, 1-2 (1995), 99–118.
- Rohan Sawhney and Keenan Crane. 2017. Boundary First Flattening. *ACM Trans. Graph.* 37, 1, Article 5 (dec 2017), 14 pages. <https://doi.org/10.1145/3132705>
- 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).
- Dmitry Sokolov. 2022. Supplemental material for “Practical lowest distortion mapping”. <https://github.com/ssloy/QIS>. Accessed: 2022-01-27.
- 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>
- Boris Springborn, Peter Schröder, and Ulrich Pinkall. 2008. Conformal Equivalence of Triangle Meshes. *ACM Trans. Graph.* 27, 3 (aug 2008), 1–11. <https://doi.org/10.1145/1360612.1360676>
- 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.