Graphical Models (2018)

Geometry-based distortion measures for space deformation

Alexander Naitsat[a,*], Emil Saucan[a,b], Yehoshua Y. Zeevi[a]

[a] *Electrical Engineering Department, Technion, Haifa, Israel*
[b] *Applied Mathematics Department, ORT Braude College, Karmiel, Israel*

ARTICLE INFO

Keywords: Digital Geometry Processing, Shape Deformation, Distortion Minimization, Quasi-Conformal Mappings, Triangle and Tetrahedral meshes.

ABSTRACT

We present a framework for optimizing a rich family of geometry-based energies defined on planes, surfaces and volumetric domains. Our approach is based on the concept of first order distortion measures and on the steepest descent optimization. Specifically, we present an algorithm for inducing optimal deformations for triangular and tetrahedral meshes. The resulting techniques apply to a variety of geometry processing problems, including ones that are highly non-convex. Among these problems are deformation and parametrization of planes, surfaces and volumetric meshes, surface mapping using volumetric textures, generation of triangular and tetrahedral meshes. In particular, the proposed techniques can be employed to devise "as close to being conformal as possible" mappings and other deformations that are nearly optimal with respect to related distortion measures, such as the isometric distortion and the distortion of a local volume. Tests, carried out on 2D and 3D data, show that the optimization process is numerically stable and fast-converging. Our approach is general and it can be run in parallel processes.

1. Introduction

A wide class of problems considered in geometry processing and computer graphics involve computation of mappings between domains in Euclidean space. These problems are often formulated in terms of geometric energies that assess the quality of a map. Let f be a smooth mapping of a domain $\mathcal{S} \in \mathbb{R}^n$ and consider the following energy

$$E(f) \triangleq \int_{\mathcal{S}} \omega(x)\,\mathrm{dist}\,(f,x)dx, \tag{1}$$

where $\omega(x)$ is a normalized cost function defined over domain $\mathcal{S}$, and $\mathrm{dist}(f,x)$ is the energy density at point x that depends on f and varies between different geometric problems. Minimizing $E(f)$ yields the highest quality mapping, and thus it provides a preferable mapping between $\mathcal{S}$ and a given target domain $\mathcal{T}$ (see Fig. 1). Formally, we consider the following fundamental problem:

Problem 1. Let $\mathcal{S}$ and $\mathcal{T}$ be compact domains with non-empty interiors in $\mathbb{R}^n$ ($n \geq 2$). Assume that $\mathcal{S}$ and $\mathcal{T}$ are homeomorphic, i.e., there is a 1:1 continuous map of one domain onto the other, and let $\mathrm{Def}(\mathcal{S},\mathcal{T})$ be a certain subset of smooth deformation functions from $\mathcal{S}$ to $\mathcal{T}$. Construct an optimal mapping

$$f_{\mathrm{opt}} = \underset{f\in\mathrm{Def}(\mathcal{S},\mathcal{T})}{\mathrm{argmin}} \int_{\mathcal{S}} \omega(x)\,\mathrm{dist}\,(f,x)\,dx, \tag{2}$$

where the fitness of a mapping is measured with respect to an energy density $\mathrm{dist}(f,x)$, called the *distortion* of f at x.

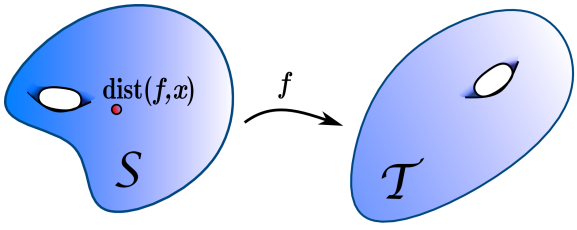

Fig. 1: An energy density $\mathrm{dist}(f,x)$ of a spatial mapping f.

A most common constraints for (2) considered in this work is the fixation of f on a subset $\mathcal{A} \subset \mathcal{S}$, i.e., a spatial constraint

$$f|_{\mathcal{A}} = g\,, \tag{3}$$

where $g : \mathcal{S} \mapsto \mathcal{T}$ is a given continuous function.

We refer to Problem 1 as *the optimal mapping problem*. In practice we are more interested in approximated solutions than in global minimizers of (2), because the proposed problem is too general to be solved under practical restrictions and

*Corresponding author: Tel.: +972-04-829-4722;
e-mail: anaitsat@campus.technion.ac.il (Alexander Naitsat)

constraints. We should, therefore, focus on a narrower class of density functions $\mathrm{dist}(f,x)$, which will be referred to as *first order distortion measures*, where "first order" refers to the fact that $\mathrm{dist}(f,x)$ does not depend on the 2^{nd} and higher order derivatives of the mapping. Despite these restrictions, the proposed set of measures is capable of quantifying a very wide range of geometric properties and it can be characterized using Jacobian singular values. For instance, various metrics of conformal distortion (such as the aspect-ratio [AL13], n-D conformal distortion [LLL15], angle and volume energies [PP12], quasi-conformal dilatation [NSZ15, NCQ*18] and least-square conformal energy [LPRM02]), isometric distortion [NSZ16, NSY17], ARAP [SA07] and Dirichlet energies [KABL14], all of which are expressed as functions of Jacobian singular values, and therefore readily fit into our optimization routine.

Although our algorithms are implemented for triangular and tet meshes, the underlined technique can be directly extended to general polyhedral meshes and to simplicial complexes embedded in $\mathbb{R}^n$. Therefore, we consider the following discrete formulation of Problem 1:

$$\underset{f:V\to\mathbb{R}^n}{\operatorname{argmin}} \sum_{v\in V} \omega(v)\,\mathrm{dist}\,(f,v)\,, \tag{4}$$

where V is the vertex set of the source mesh that represents $\mathcal{S}$, $\omega(v)$ is a vertex weight and $\mathrm{dist}(f,v)$ is an approximation of the continuous distortion, $\mathrm{dist}(f,x)$, in the vicinity of a vertex v. The spatial constrains defined by (3) are expressed in the discrete case as

$$f(v_i) = f_i \text{ for } v_i \in V,\ i = 1,\ldots,m.$$

Even for a simple distortion metric $\mathrm{dist}(f,x)$, the exact solution of Problem 1 may give rise to a fundamental difficulty, since in the general case there is very little known about properties of global minimizers, or even about their existence. For example, choose a measure of the *conformal distortion* that assesses how far is f from being conformal in a neighborhood of a point x. In this case, the global minimizers are closely related to the so-called *extremal quasi-conformal mappings*. To the best of our knowledge, there does not exist a technique to characterize, even qualitatively, the set of extremal quasi-conformal mappings for dimensions $n \geq 3$. The latest research in this area provides only a moderate lower bounds of conformal distortion for a basic source and target domains, such as convex polyhedrons, n-balls and wedges [Car74, V̈71]. Therefore, our framework is not intended to provide an exact solution for the optimal mapping problem, instead we consider approximate solutions of (2). In particular, for a given mapping $f : \mathcal{S} \to \mathcal{T}$ we are looking for techniques that can iteratively refine the mapping till the process reaches an admissible degree of convergence.

Our algorithm is based on a modified gradient descent that does not require any constraints on the optimized energies, except the "first order" property and trivial geometric assumptions. Compared with other techniques, our method is more general and significantly simpler in the implementation, while it is competitive in terms of the performance and accuracy. Particularly, on noisy meshes we achieve results that compare favorably to state-of-art methods. Moreover, unlike other more sophisticated techniques, our algorithms can be naturally parallelized using simple mesh cutting.

1.1. Related work

There are plenty of techniques for construction of planar deformations via minimization of geometric energies. However, most of these methods employ mathematical tools (e.g., complex analysis) that cannot be extended beyond 2D. Therefore, we focus on related techniques that operate on volumetric domains or general studies that deal with arbitrary dimensions $n \geq 2$. The relevant methods can be qualitatively divided into the following groups:

Convexification methods (CM). These include projecting simplicial mappings on the space of bounded distortion maps (BD and LBD) [AL13, KABL15] and controlling singular values via semi-definite programming (SDP) [KABL14]. In general, these techniques approximate problems similar to (2) by a sequence of nested convex problems for which convex optimization tools can be applied. Similarly to our technique, CM methods express geometric energies as functions of the Jacobian of a simplicial mapping. However, the use of convex optimization tools imposes certain constraints on CM. These techniques are therefore not as generic as our approach. Notably, our research deals with the general notation of the spatial deformation that extends beyond $\mathbb{R}^3$, while the CM techniques, operate only on a narrow finite subset of convex and quasi-convex measures. Moreover, in most cases CM methods deal only with orientation-preserving maps and assume monotonicity with respect to singular values, which restrict these technique to even smaller subsets of geometric measures and deformations.

Harmonic mappings (HM). Minimization of 3D harmonic energy under a fixed boundary constraints yields to the discrete harmonic mappings [WGY*03, LGW*07]. These methods are employed in [CPS15] to minimize quasi-conformal errors for volume deformations. The method of [BCWG09] considers linear combinations of the harmonic basis functions that minimize the rigidity energy. HM techniques produce smooth deformations over a wide range of spatial domains. However, compared with our approach, HM methods operate on a restricted space of harmonic functions and consider a limited class of associated geometric measures. Notably, our algorithm is capable of minimizing a linear combination of various measures, including the smoothness or Dirichlet energy which is closely related to HM.

Most-Isometric parametrization (MIPS). The classical MIPS energy [Hor01] is extended in [FLG15] to advanced MIPS energies optimized over triangular and tetrahedral domains by indirect block coordinate descend. However, [FLG15] and other related MIPS techniques are restricted to few conformal and isometry-based energies, and they do not consider a general framework of generic distortion measures and their properties. Similar global approaches to mesh parametrization and surface mapping are considered in BPFB [SS15] and [APL14], respectively.

Least squares projection. Least square fitting, used for construction of 2D conformal mapping, is extended in [PP12] to volumes. The method projects a volumetric mapping into coordinate planes of 3D space, and solves the corresponding system of Cauchy-Riemann equations [Ahl66, pp. 24-25]. Since

Cauchy-Riemann equation characterizes 2D conformal functions, projections of resultant volumetric maps approaches conformality. The study of [BDS*12] unifies major least squares geometry processing techniques into a framework of shape projection operators. Although shape operators can represent a variety of geometric properties, the implementation of this method requires separate algorithms for different geometric problems. In contrast, we present a single algorithm capable of minimizing the full range of spatial distortions.

Proxy and preconditioning methods. Certain geometric energies can be minimized indirectly by modifying a descend direction or by employing simpler proxy measures. For example, AQP [KGL16] uses quadratic proxy function whose Hessian is chosen to be the Laplacian. A weighted proxy terms is employed in SLIM [RPPSH17] to find a global descend direction, while AKVF [CBSS17] employs a preconditioning technique derived from Killing vector fields [SBCBG11]. Although these methods are scalable to large datasets, they consider only certain classes of problems; e.g., AKVF considers isometric distortions for codimension-zero target meshes, SLIM minimizes isometric and certain conformal energies.

Related problems. The mapping problem can be expresses in other forms than do not evolve direct optimization of (2). Among them are discrete versions of the classical mass transportation problem employed in [SDGT12], [SCL*17] and [SCQ*16] for hexahedral, tetrahedral and triangular meshes, respectively. Some methods operate on metric domains, e.g., using a circle parking metric to model Ricci flow on discrete surfaces [ZGZ*14]. Discrete Ricci flow achieves the targeted metric conformal to the original by obtaining an user-defined distribution of Gaussian curvatures at vertices. The method of [ZLYG09] computes Beltrami coefficients to obtain quasi-conformal map with given eccentricity. Since this approach is based on solving Beltrami equation, it does not extends directly to volumes. Despite some similarities, the above problems are formulated in a different manner, and thus are not compared with our algorithm in this work. For example, mass transportation methods usually operate over the space of volume preserving maps and their aim is to minimize a transportation cost function. While our method is capable of approaching volume preserving maps, it does not support transportation cost objective functions, since they are not invariant to translation and rotations. Nevertheless, the above methods can be integrated into our framework in the initialization stage.

2. Distortion measures

Generally, we consider two function spaces for the continuous Problem 1: locally injective continuous and locally injective smooth mappings between domains in $\mathbb{R}^n$, shortly referred to as *deformation* and *smooth deformation* functions, respectively. [For formal mathematical definitions see Appendix A.]

Since continuous mappings are more general, we first define a general concept of the distortion measures for deformation functions which are not necessary differentiable. The concept is based on basic geometric assumptions that include:

- According to (1), distortions are densities of geometric energies, therefore they are local measures of a mapping.
- If a mapping f is differentiable, then distortions of f are independent of its higher order partial derivatives.
- Geometric energies are characterized by changes in intrinsic geometry under a mapping. Therefore, distortions are invariant to a particular choice of orthogonal coordinates.

We formalize the above observations as follows:

Definition 2.1. A quantity $\mathrm{dist}(f,x)$, for a deformation function f and a point x in its domain, is called *a first order distortion measure of f at x*, or for short *a distortion*, if the following conditions are satisfied:

1. **First order precision:** Assume that f,g are deformation functions defined over a common domain $\mathcal{S}$. If $f(x)-g(x)=o(\|x-x_0\|)$ for $x_0\in\mathrm{int}(\mathcal{S})$, then
$$\mathrm{dist}(f,x_0)=\mathrm{dist}(g,x_0).$$
2. **Invariance to compositions with rigid transformations:** If f is a deformation function defined in $\mathcal{S}$ and T_1,T_2 are rigid coordinate transformations (i.e., compositions of translations and unitary linear operators), then for each $y\in T_1(\mathcal{S})$
$$\mathrm{dist}(T_2\circ f\circ T_1^{-1},y)=\mathrm{dist}(f,T_1^{-1}(y)).$$
In other words, the position and orientation of coordinate frames (of both the domain and the image) have no impact on distortions measured at the respected points.

Definition 2.1 is, in turn, an algebraic description of very large set of geometric measures and it is particularity useful for theoretical analysis of various problems. In practice, however, dealing with smooth deformations and their discrete approximations requires a more compact formulation of distortions. Notably, according to the following theorem, distortion measures restricted to smooth deformations can be equivalently characterized by the Jacobian singular values:

Theorem 2.1. *A first order distortion measure,* $\mathrm{dist}(f,x)$*, can be expressed as a function of singular values of the Jacobian df_x, i.e.,*

$$\exists\mathcal{D}:\mathbb{U}^n\to\mathbb{R},\ \mathrm{dist}(f,x)=\mathcal{D}\left(\sigma_1(df_x),...,\sigma_n(df_x)\right),$$

where $\mathbb{U}^n$ is the half space below the main diagonal of $\mathbb{R}^n$, namely

$$\mathbb{U}^n\triangleq\{x\in\mathbb{R}^n|x_1\geq x_2\geq\ldots\geq x_n>0\}. \tag{5}$$

[For proof see Appendix B.]

Restricting distortion measures to smooth deformations immediately implies the opposite version of Theorem 2.1.

Theorem 2.2. *Let $\mathcal{D}:\mathbb{U}^n\to\mathbb{R}$ be a function. Then,*

$$\mathrm{dist}(f,x)=\mathcal{D}\left(\sigma_1(df_x),..,\sigma_n(df_x)\right)$$

is a first order distortion measure for smooth deformations.

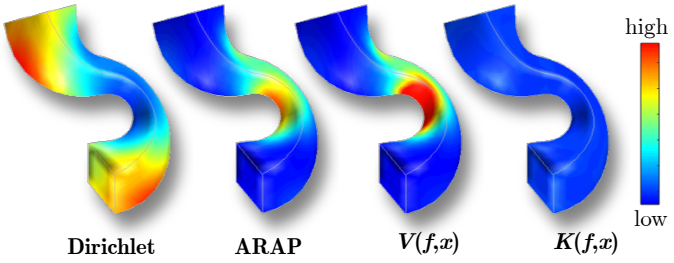

Fig. 2: Distortion measures visualized for the bar deformation.

Consequently, according to Theorems 2.1 and 2.2, we can uniquely identify each n-dimensional distortion $\mathrm{dist}(f,x)$ for smooth deformation f with a scalar function $\mathcal{D} : \mathbb{U}^n \to \mathbb{R}$. We call the latter *the canonical representation* of a measure $\mathrm{dist}(\cdot,\cdot)$.

2.1. Examples of distortion measures

Following the definition, distortion values of a map f can be intuitively interpreted as measures of f's rigidity. Rigid transformations form the space of length-preserving (isometric) maps which, in turn, is the intersection of two sets containing angle-preserving (conformal) and volume-preserving (equi-volume) maps. Thus, distortions incorporate the following major groups: length-wise measures, angle-wise measures and measures based on assessments of local volumes. In view of the above observations, we recall the differential definitions of isometric, conformal and volume-preserving mappings, formulated for smooth deformations as follows:

A smooth deformation f is a *conformal mapping* of a domain $\mathcal{S} \subset \mathbb{R}^n$ if for each point $x \in \mathrm{int}(\mathcal{S})$ it scales the space uniformly in every direction. This can be stated formally as

$$\|df_x \cdot u_1\| = \|df_x \cdot u_2\|\,, \tag{6}$$

where df_x denotes the Jacobian matrix at a point x of function f and $u_1, u_2 \in \mathbb{R}^n$ are arbitrary unit vectors. Following this notation, a conformal map $f(x)$ is *isometric* in $\mathcal{S}$ if

$$\forall x \in \mathrm{int}(\mathcal{S}) : |\det df_x| = 1\,, \tag{7}$$

where (7) alone defines a class of *volume-preserving*, or *equi-volume* transformations.

While there is an abundance of conformal mappings in 2D, general higher dimensional domains can be mapped only *quasi-conformally*, which means that condition (6) is satisfied only partially. In this case we say that the mapping produces *conformal distortion* of the space. Applying the same terminology for isometries and equi-volume maps, yields the concepts of *isometric* and *volume distortions*, respectively. These quantities are local properties associated with a spatial deformation between given source and target regions. As shown by Theorem 2.1, we can express distortions of Euclidean space under a mapping f as a simple function of singular values of Jacobian

$$\sigma_i = \sigma_i(df_x) \text{ for } i = 1,\ldots,n.$$

Since the equivalent condition for a smooth f to be conformal is

$$\sigma_i = \sigma_j, \text{ for } 1 \le i,j \le n\,,$$

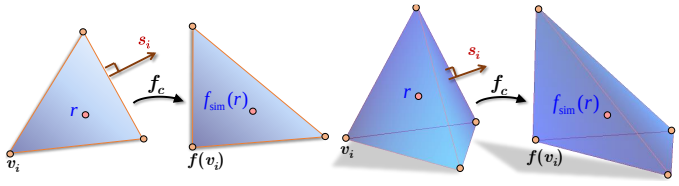

Fig. 3: Simplicial mapping of a cell and the corresponding constructions.

conformal distortions are, in fact, estimates of the variation of singular values and they are often measured by the following quantities:

- *Quasi-conformal (qc) dilatation*

$$K(f,x) \triangleq \max\left\{\frac{\sigma_1\cdots\sigma_{n-1}}{{\sigma_n}^{n-1}}, \frac{\sigma_1^{n-1}}{\sigma_2\cdots\sigma_n}\right\},$$

often defined as

$$K(f,x) = \max\{K_I(f,x), K_O(f,x)\}\,, \tag{8}$$

where K_I and K_O are the so-called *inner* and *outer* qc-dilatations. These quantities can be visualized as volume ratio between a small ellipsoid, obtained by mapping an infinitesimal sphere under f, and its inscribed and circumscribed balls.

Quasi-conformal dilatations are often employed in mathematical analysis of *qc*-mappings; notably they are useful in estimations of geometry-dependent bounds of conformal distortions.

- *Condition number*, also called *linear dilatation*

$$\kappa(f,x) \triangleq \frac{\sigma_1}{\sigma_n}\,, \tag{9}$$

which is the ratio between the maximal and the minimal singular values. Appears as a metric of conformal distortion in [AL13, KABL14] for dimensions $n = 2,3$ ($\kappa(f,x) = K(f,x)$ in 2D). This quantity is often employed in numerical analysis and other disciplines outside the geometry scope.

- *MIPS energy*, defined in 2D as

$$\mathrm{MIPS}_{2\mathrm{D}}(f,x) \triangleq \frac{\sigma_1}{\sigma_2} + \frac{\sigma_2}{\sigma_1} = \frac{\sigma_1^2+\sigma_2^2}{\sigma_1\sigma_2}\,. \tag{10}$$

Although historically MIPS is refered to as "most isometric parametrizations", this energy measures how σ_1 differs from σ_2, and thus it is a metric of conformal distortion. The energy were extensively employed in early mesh processing applications, since for the discrete problem (4) $\mathrm{MIPS}_{2\mathrm{D}}(f,v)$ is proven to be locally convex as the function of a single vertex image $f(v)$ [Hor01] . MIPS can be extended to n-dimensions as n-times the ratio between arithmetic and geometric means of $\sigma_1^2,\ldots,\sigma_n^2$

$$\mathrm{MIPS}_{n\mathrm{D}} \triangleq \frac{\sigma_1^2+\cdots+\sigma_n^2}{(\sigma_1\cdots\sigma_n)^{2/n}} = \frac{\mathrm{trace}(df_x^T df_x)}{|\det(df_x)|^{2/n}}\,. \tag{11}$$

We now consider metrics of the isometric distortion which are direct measures of the rigidity. Since singular values of an isometry equal 1, these distortions assess the deviation of $(\sigma_1,...,\sigma_n)$ from the vector $(1,...,1)$, most commonly, using the following quantities:

- *Quasi-isometric (qi) dilatation*

$$C(f,x) \triangleq \max\{\sigma_1, \sigma_n^{-1}\}, \tag{12}$$

where σ_1 and σ_n are, actually, local measures of the expansion and the contraction of the space, respectively. This measure is used in both theoretical studies [Car74] and in practice [NSZ16, KABL15], where it is often normalized according to the relative sizes of the source and target domains.

- *Rigidity energy*, also often called ARAP (as rigid as possible) energy, was initially introduced for 2D surfaces [SA07] and its direct extension to $\mathbb{R}^n$ is defined by

$$E_{\mathrm{ARAP}}(f,x) \triangleq \sum_{i=1}^{n} (\sigma_i - 1)^2 . \tag{13}$$

- Symmetric rigid energy, also called *symmetric Dirichlet energy* (see the discussion in Appendix C) can be defined in $\mathbb{R}^n$ by

$$E_{\mathrm{SD}}(f,x) \triangleq \frac{1}{2n}\sum_{i=1}^{n} \left(\sigma_i{}^2 + \sigma_i{}^{-2}\right) . \tag{14}$$

Finally, if f is equi-volume in the vicinity of x, then

$$|\det(df_x)| = \sigma_1\sigma_2\cdots\sigma_n = 1 .$$

Hence, for estimation of the volume distortion (area distortion in 2D) we employ the following measure

$$V(f,x) \triangleq \max\left\{|\det(df_x)|, |\det(df_x)|^{-1}\right\}, \tag{15}$$

where $|\det(df_x)|$ and $|\det(df_x)|^{-1}$ can be considered as assessment of the dilatation and compression of a local volume, respectively.

Another measure widely used in applied geometry is *Dirichlet energy*. This energy, also referred to as the *smoothness energy*, is traditionally employed in construction of harmonic surface parametrization. Few approaches extend it to volumetric domains [WGY*03, LGW*07]. We suggest the following direct extension of the Dirichlet energy density to n-dimensions:

$$E_{\mathrm{Dirichlet}}(f,x) = \frac{1}{n}\sum_{i=1}^{n} \sigma_i{}^2 . \tag{16}$$

Obviously, the above quantities are formulated for the n-dimensional case, where in dealing with actual planar and volumetric distortions we take $n=3$ and $n=2$, respectively. The above distortion metrics are visualized in Fig. 2 using a simple example of deforming rectangular bar in 3D.

Note that minimizing different metrics of the same distortion may lead to distinct results as illustrated in Fig. 12 for $C(f,x)$ and $E_{\mathrm{ARAP}}(f,x)$ measures. Although both measures represent isometric distortion, $C(f,x)$ depends only on σ_1 and σ_n, while ARAP energy is a function of all singular values, and thus it considers n independent directions of the space.

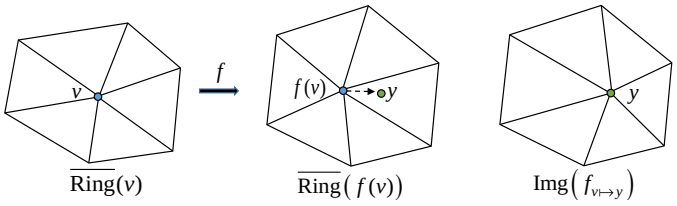

Fig. 4: Construction of the mapping $f_{v\mapsto y}$ defined according to (26) in the ring of the neighboring vertices of v (a 2D case is illustrated).

Note that the deformation function space, and thus the entire Problem 1, are highly non-convex, since for any deformation f in the space the line $(1-\lambda)f+\lambda(-f)$ contains a singular map. Since simplifying the original problem and restricting geometric measures are beyond our goals here, the above discussion does not analyze which distortions are convex. For such analyses, the reader may refer to convexification or preconditioning methods where distortion convexity plays a substantial role. For example, consider an approximation of extremal qc-mappings. In this case, SDP employs the condition number because it is a quasi-convex function of the singular values, while SLIM is applicable for minimizing the MIPS energy, since it is locally convex with respect to a single target vertex in 2D.

2.2. Simplicial mappings

First, we consider a discrete representation of a n-dimensional deformation $f : \mathcal{S} \to \mathcal{T}$, denoted by $f = \left(f^{(1)},...,f^{(n)}\right)$. We represent the source domain as a *simplicial mesh* $(\mathcal{V},\mathcal{C})$, where $\mathcal{V}$ is the set of vertices and $\mathcal{C}$ is a set of oriented simplices, called *cells*, that constitute a triangulation of $\mathcal{S}$. We use general formulations for dimension $n \geq 2$, where a cell is a n-simplex which in geometric terms is a convex hull of $n+1$ vertices in $\mathbb{R}^n$.

We approximate f on a cell $c = (v_1, v_2, ..., v_{n+1})$ by the affine function f_c that satisfies $f_c(v_i) = f(v_i)$ for each i. Then, a piecewise affine function

$$f_{\mathrm{sim}} : \bigcup_{c\in\mathcal{C}} c \to \mathbb{R}^n, \;\; f_{\mathrm{sim}}(x) = f_c(x) \text{ for } x \in c ,$$

is called *a simplicial map*. It can be obtained in cell c as a linear blending of the samples $f(v_1),...,f(v_{n+1})$ as follows: if r is a point inside c, then

$$f_{\mathrm{sim}}(r) = \lambda_1 f(v_1) + \lambda_2 f(v_2) + ... + \lambda_{n+1} f(v_{n+1}) , \tag{17}$$

where $\lambda_1,..,\lambda_{n+1}$ are barycentric coordinates of r, i.e.,

$$\sum_{j=1}^{n+1} \lambda_i = 1 \text{ and } \sum_{j=1}^{n+1} \lambda_i v_i = r . \tag{18}$$

Practically we are interested in 2D and 3D cases, where $\mathcal{C}$ is a set of triangles and tetrahedrons, respectively. In these dimensions we use the following constructions (Fig. 3)

- In 2D: Let μ_i be the edge of a triangle c against v_i, and s_i be a vector along the outwards facing normal of μ_i, such that $\|s_i\| = \text{length}(\mu_i)$.
- In 3D: Let μ_i be the face of tetrahedron c against v_i, and s_i be a vector along the outwards facing normal of μ_i such that $\|s_i\| = \text{area}(\mu_i)$.

Then, for $n = 2, 3$ the value of a simplicial map at $r \in c$ is

$$f_c^{(i)}(r) = \frac{1}{n\,\text{volume}(c)} \sum_{j=1}^{n+1} (v_{n+1} - r) \cdot s_j f^{(i)}(v_j)\,, \tag{19}$$

and rows of the Jacobian df_c are gradients

$$\nabla f_c^{(i)} = \frac{-1}{n\,\text{volume}(c)} \sum_{j=1}^{n+1} s_j f^{(i)}(v_j),\ \ i = 1, \ldots, n\,. \tag{20}$$

This scheme for Jacobian computation can be obtained equivalently by solving linear systems (17) and (18), or by computing gradients of the, so called, *hat functions* that constitute the basis of the simplicial mapping space.

Finally, choose a distortion measure and express it as a function $\mathcal{D}$ of the Jacobian singular values. Then, in the cell c the distortion has a constant value

$$\text{dist}(f_c) = \mathcal{D}\left(\sigma_1(df_c), \ldots, \sigma_n(df_c)\right). \tag{21}$$

The distortion at $v \in \mathcal{V}$ can be therefore estimated by the weighted average

$$\underset{w}{\text{dist}}(f_{\text{sim}}, v) = \sum_{c \in \text{Cell}(v)} w(c, v)\,\text{dist}(f_c)\,, \tag{22}$$

where $\text{Cell}(v)$ is a set of the neighboring cells sharing v and $w(c, v)$ are positive weights. The most widely employed weights are the *uniform* $w(c, v) \equiv 1$, *volume weights* $w(c, v) = \text{volume}(c)$ and the *average volume weights*:

$$w(c, v) = |c|^{-1}\,\text{volume}(c)\,, \tag{23}$$

where $|c|$ denotes the number of vertices in c.

The above numeric scheme for planar deformations can be employed directly for simplicial maps between 3D triangular meshes (i.e., triangular meshes embedded in $\mathbb{R}^3$). The only difference that should be considered in (19) is that on surfaces s_i are defined on a general plane containing a 3D triangle c. In this case $\sigma_3(df_c) = 0$, and therefore we set $\text{dist}(f_c) = \mathcal{D}(\sigma_1, \sigma_2)$, where $\mathcal{D}$ is a given 2D distortion metric.

3. Minimizing distortion measures

We look for simplicial mappings that approximate solutions of (2). Let us reexamine (4) by employing the notion of simplicial map f of a mesh $(\mathcal{V}, \mathcal{C})$, the concept of distortion measure $\text{dist}(f, x)$ and its estimates on vertices defined by (22). Then, the discrete energy involved in (4) can be rewritten as follows:

$$E(f) = \sum_{v \in V} \omega(v) \underset{w}{\text{dist}}(f, v) \tag{24}$$

$$= \sum_{c \in \mathcal{C}} w(c)\,\text{dist}(f_c)\,, \tag{25}$$

Algorithm 1: Minimization of distortion measures

Input:

- Source mesh $(\mathcal{V}, \mathcal{C})$ and set of fixed vertices $\mathcal{V}_{\text{fix}}$.
- Initial simplicial map f_0.
- Distortion measure $\text{dist}(f, x)$ for the minimization.
- Cell and vertex weights: $\omega(v)$ and $w(c, v)$.
- Maximal number of iterations N and convergence threshold ε.

1. $\text{dist}(f_{-1}) \leftarrow \infty$, $i \leftarrow 0$.
2. $W(c) \leftarrow \sum_{v \in c} \omega(v) w(c, v)$ for each $c \in \mathcal{C}$.
3. $\text{dist}(f_0) \leftarrow \sum_{c \in \mathcal{C}} W(c)\,\text{dist}((f_0)_c)$.

while $i \le N$ *and* $\text{dist}(f_{i-1}) - \text{dist}(f_i) > \varepsilon$ **do**

$f_{i+1} \leftarrow \text{GreedyGradDescent}(\mathcal{V}, \mathcal{V}_{\text{fix}}, \mathcal{C}, \text{dist}, f_i, \omega, w)$;
$\text{dist}(f_{i+1}) \leftarrow \sum_{c \in \mathcal{C}} W(c)\,\text{dist}(f_{i+1}, c)$;
$i \leftarrow i + 1$;

Output: Simplicial map f_i obtained at the last iteration

where, according to (22), we set[1]

$$w(c) \triangleq \sum_{v \in c} \omega(v) w(c, v)\,.$$

Eventually, the optimal mapping problem is reduced to minimizing $E(f)$, equivalently defined by (25) and (24), under the constraints

$$f(\mathcal{V}_{\text{fix}}) = f_0(\mathcal{V}_{\text{fix}})\,,$$

where $\mathcal{V}_{\text{fix}} \subset \mathcal{V}$ is the set of vertices where f is fixed and f_0 is a given initial deformation.

To resolve the problem, we consider a local approach based on the fact that a displacement of a vertex v inside the hull of its neighboring cells does not affect values of $\text{dist}(f, u)$ if u does not share a common cell with v. Denote by $\text{Ring}(v)$ the set of the neighboring vertices of v, called *the ring of* v, and denote by $\overline{\text{Ring}}(v)$ the hull of the ring, namely

$$\overline{\text{Ring}}(v) = \bigcup_{c \in \text{Cell}(v)} c.$$

Then, we solve the following formal problem for a vertex v

$$\underset{y \in \overline{\text{Ring}}(f(v))}{\text{argmin}}\ \underset{w}{\text{dist}}(f_{v \mapsto y}, v), \tag{26}$$

where $f_{v \mapsto y}$ denotes the simplicial map obtained by restricting f to fixed values on $\text{Ring}(v)$ and setting $f_{v \mapsto y}(v) = y$ (see Fig. 4). Our aim is to minimize $E(f)$ locally by solving (26) for each $v \in \mathcal{V} \setminus \mathcal{V}_{\text{fix}}$. Practically, we prefer a greedy approach to

[1] Taking average volume weights and $\omega(v) \equiv 1$ implies $w(c) = \text{volume}(c)$.

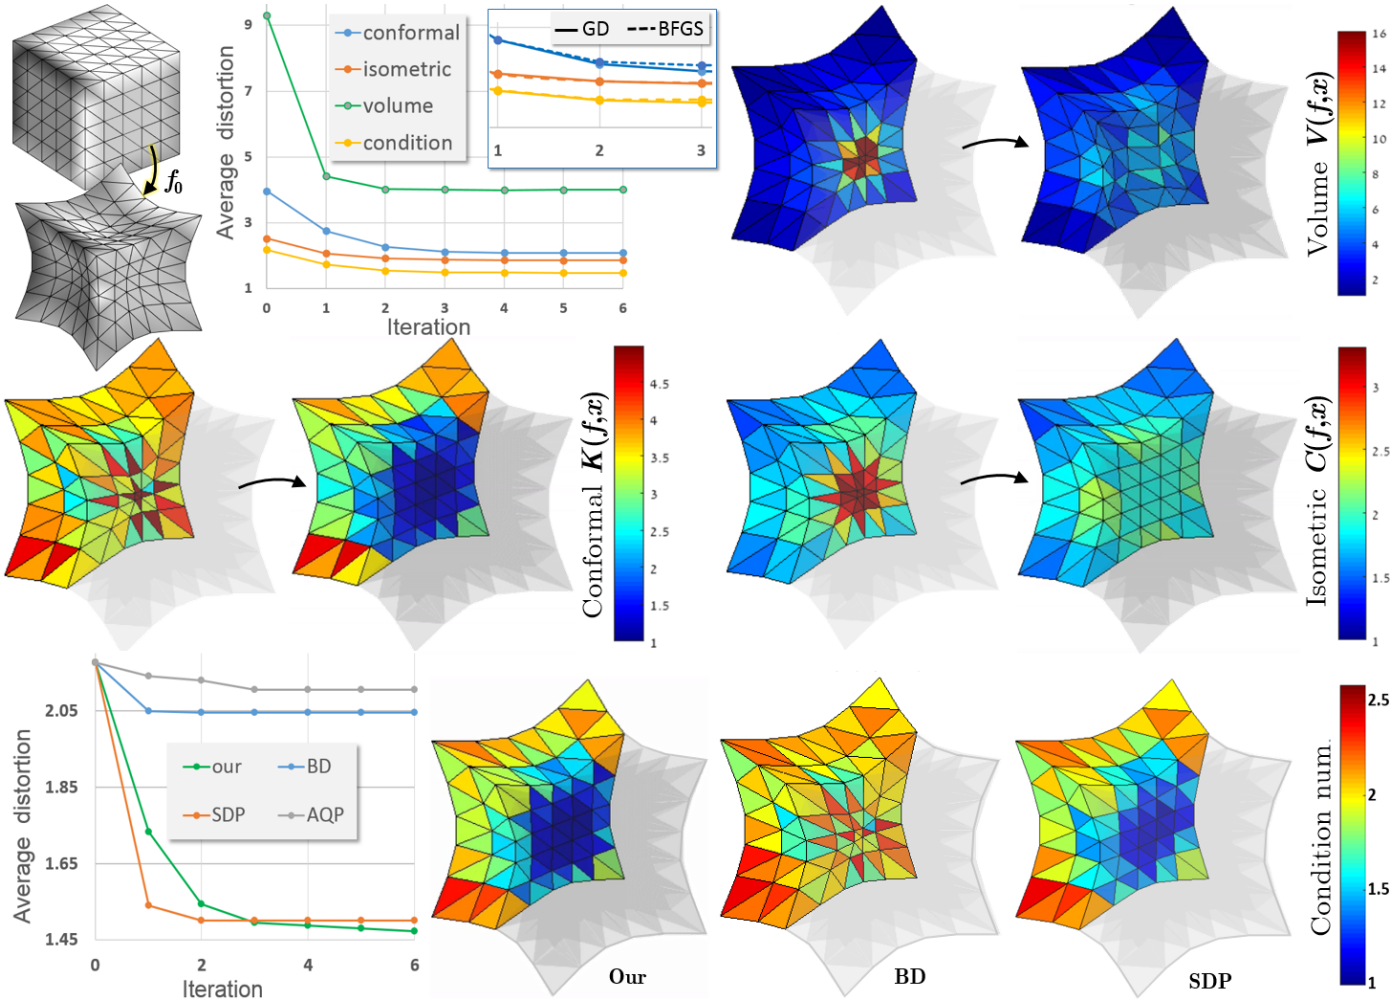

Fig. 5: Minimizing distortions under fixed boundary constraints. The initial map f_0 is shown on top left. First two rows illustrate result of Algorithm 1 for different distortions, including comparison between GD and BFGS solvers for iterations 1-3. Arrows point from the image of f_0 to the image of the optimal deformation obtained after 6 iterations. Cross sections are colored according to distortions per tetrahedron. The third row shows the comparison of our algorithm with related methods. Note that AQP method was employed indirectly via minimization of isometric distortion and that LBD method failed to reduce condition number while preserving the boundary constraints.

the problem; i.e., we optimize at each step vertex $v \in \mathcal{V} \setminus \mathcal{V}_{\text{fix}}$ that achieves the maximal distortion

$$\max_{v \in \mathcal{V} \setminus \mathcal{V}_{\text{fix}}} \omega(v) \operatorname*{dist}_{w}(f, v).$$

Then, we fix the position of $f(v)$ to the attained solution of (26) and continue to the next ring, until f is fixed on all vertices.

Our implementation supports GD (gradient descent) and BFGS [WN99] methods for solving (26). Both methods achieve very similar numerical results, while BFGS is more computationally expensive [2]. The procedure for gradient descent is presented in Algorithm 2, where the notion $\nabla \operatorname{dist}(f, v, y)|_{y=f(v)}$ refers to the gradient of the function $\operatorname{dist}(f_{v \mapsto y}, v)$ computed at $y = f(v)$. The whole procedure is repeated by Algorithm 1 till the energy $E(f_i)$ converges or till the maximal number of iterations is reached. Eventually, the output of the whole process is a simplicial map f_N that depends on the following input parameters: initial simplicial map f_0 of the mesh $(\mathcal{V}, \mathcal{C})$; canonical representation of distortion measures $\mathcal{D}(\sigma_1, ..., \sigma_n)$; spatial constraints expressed by the set of the fixed vertices $\mathcal{V}_{\text{fix}}$; vertex weights $\omega : \mathcal{V} \to (0, \infty)$ indicating the "distortion costs" over the target domain and the cell weight $w(c, v)$ used in the estimation of distortions on vertices.

Further, we present a number of applications based on Algorithm 1 that include: deformation of meshes, mesh generation, improving volumetric parametrization and mapping surfaces with complex topology structures. The quantitative results of our applications are depicted using average distortion values over the volume of a source domain. To avoid a long list of parameters employed for each result, we assume uniform vertex weights (i.e, $\omega(v) \equiv 1$), average volume cell weights (23) and the use of the gradient descent solver (Algorithm 2), unless stated otherwise.

The complexity of each iteration of Algorithm 1 is $m \log(m)$ with respect to the number of vertices, and the total number of iterations depends on geometric complexity. For instance,

[2] Similarly, the Newton method is not efficient in our case; it requires heavier numerical computations than GD, and it cannot boost the convergence due to the highly non-convex structure of the general problem.

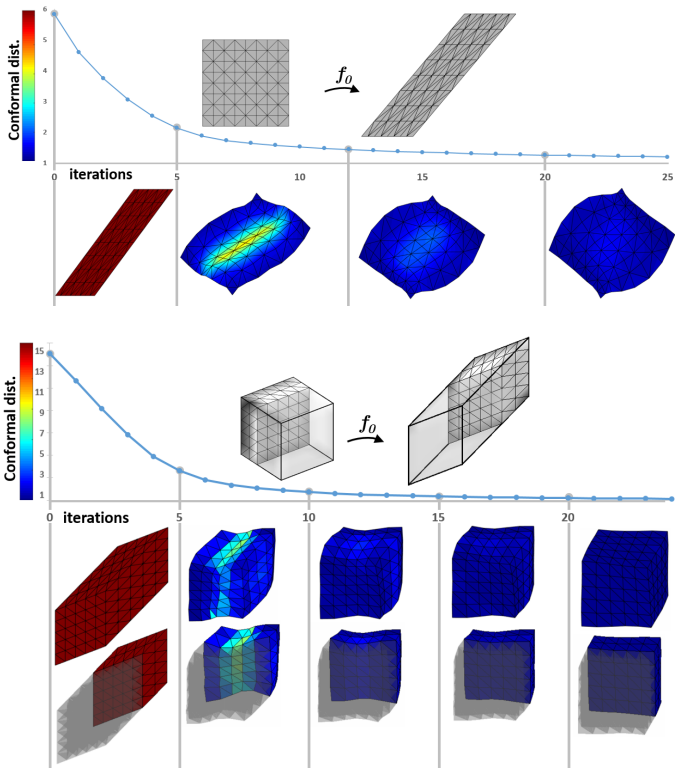

Fig. 6: Unconstrained minimization of conformal distortion $K(f,x)$ in 2D (top) and in 3D (bottom). Initial shear mapping of square and cube are shown above plots, and images of deformations f_i for chosen iterations are depicted below.

as illustrated in Fig. 5, if f_0 is injective, then the optimization process converges after a few iterations. However, as seen in Fig. 10, the presence of self intersections and nearly singular cells slows down the process in 3D. Note that according to Fig. 8 this impact on the performance is not significant in 2D.

Since using analytic SVD computations for 3×3 matrices is impractical, the general implementation of Algorithm 2 computes distortion gradients numerically using the Forward-Backward Euler scheme.

Fig. 6 illustrates the process of minimizing conformal distortion without constraints (i.e, $\mathcal{V}_{\text{fix}}=\emptyset$) in 2D and in 3D. Apparently, the resultant process is a *conformal flow* of $\text{Im}(f_0)$,i.e., a movement of vertices toward the image of the nearest conformal map. We chose a simple case, where source domains are a square and a cube, and initial deformations are shear transformations. As Fig. 6 illustrates, our algorithm in 3D leads to a rotation and a slight uniform scaling of the input cube, wheres in 2D the input square is slowly deformed into a different shape. These results emphasize the fundamental difference between the continuous plane, where any pair of simply-contained domains can be mapped conformaly onto each other, and higher dimensions, where conformal mapping are restricted to the limited set of Möbious transformations.

3.1. Parallelization

One of major advantages of our local approach is that Algorithm 1 can be naturally parallelized by dividing source domain into sub regions. Let denote by a 4-tuple $\mathcal{P}=(f_0,\mathcal{V},\mathcal{C},\mathcal{V}_{\text{fix}})$ a

Algorithm 2:
GreedyGradDescent$(\mathcal{V},\mathcal{V}_0,\mathcal{C},\text{dist},f_0,\omega,w)$

Input:

- Source mesh $(\mathcal{V},\mathcal{C})$ and initial map f_0 fixed on $\mathcal{V}_0$.
- Distortion $\text{dist}(f,x)$ in the canonical representation.
- Vertex weights $\omega(v)$ and cell weights $w(c,v)$.

Initialize: $\mathcal{V}_{fix}\leftarrow\mathcal{V}_0$, $f\leftarrow f_0$.

while $\mathcal{V}\setminus\mathcal{V}_{fix}\neq\emptyset$ **do**
  $v_{\max}\leftarrow \underset{v\in\mathcal{V}\setminus\mathcal{V}_{fix}}{\text{argmax}}\ \omega(v)\underset{w}{\text{dist}}(f,v)$;
  repeat
    $\nabla\leftarrow\nabla\ \underset{w}{\text{dist}}(f,v_{\max},y)\Big|_{y=f(v_{\max})}$;
    Using line search choose Δt such that
    $(f(v_{\max})-\Delta t\nabla)\in\overline{\text{Ring}}\,(f(v_{\max}))$;
    Refine f at $v_{\max}$: $f(v_{\max})\leftarrow f(v_{\max})-\Delta t\nabla$;
  until *Convergence of* $\text{dist}(f,v_{max})$;
  $\mathcal{V}_{fix}\leftarrow\mathcal{V}_{fix}\cup\{v_{\max}\}$;

Output: Simplicial map f obtained by optimizing f_0.

problem of minimizing $\text{dist}(f,x)$ on a mesh $(\mathcal{V},\mathcal{C})$ under fixed vertices $\mathcal{V}_{\text{fix}}$ and inital map f_0. The problem can be processed in parallel as series of sub-problems:

$$\mathcal{P}_i=(f_{0,i},\mathcal{V}_i,\mathcal{C}_i,\mathcal{V}_{\text{fix}}\cap\mathcal{V}_i)\,,\ i=1,...,k, \tag{27}$$

where $(\mathcal{V}_i,\mathcal{C}_i)$ are sub-meshes with disjoint cells that constitute a partition of $(\mathcal{V},\mathcal{C})$ and $f_{0,i}$ denotes the restriction of f_0 to $\mathcal{V}_i$.

The parallel optimization is valid as long as the simplicial mappings for each $\mathcal{P}_i$ coincide on their common vertices $\mathcal{V}_{\text{com}}$ (these that belong to at least two sub-meshes). Our approach, presented in Algorithm 3, ensures the validity by switching sequentially between optimization of $\mathcal{V}_{\text{com}}$ and other vertices. Each iteration in Algorithm 3 first optimizes $(\mathcal{V}_\partial,\mathcal{C}_\partial)$ mesh that covers common vertices. Then, $\mathcal{V}_\partial$ are fixed and problems of (27) are solved in parallel via Algorithm 2 (for short we dropped some input parameters of Algorithm 2, including cell and vertex weights). In order to use efficiently available resources for parallel computations it is necessary to satisfy $|\mathcal{V}_\partial|=o(\max|\mathcal{V}_i|)$. If this criterion is not met, then optimization of $(\mathcal{V}_\partial,\mathcal{C}_\partial)$ can be run in parallel via the recursive call of Algorithm 3. Note that mapping update steps in Algorithm 3 are not necessary if the mesh slice dataset is shared between multiple threads.

Tests of Algorithm 3 were carried out using a spatial partition of the source domain with varying resolution and number of sub-meshes (k). The results are shown in Table 1 for 10^{-4} convergence threshold. For high resolution models the best results we achieved with k equals the number of available CPU cores. However, in low-to-middle resolution models, taking k as a half core number attains the best results due to the multi-thread initialization cost and un-parallel optimization of $\mathcal{V}_{\text{com}}$ vertices. According to the results, even the single-thread Algorithm 1 achieves superior performance compared with related methods which do not support parallel computing.

3.2. Stability and injectivity

The gradient descent allows minimization of arbitrary distortions. However, to ensure stability and injectivity one should consider a certain subset of measures, referred to as *regular distortion measures* (RDM).

Suppose that we deal only with smooth bijective maps. Then, $\mathrm{dist}(f,x)$ is *normalized* if it is minimal when f is rigid in a neighborhood of a domain point x and the distortion is infinite if df_x is singular or infinite, except when all singular values approach 0 or ∞ simultaneously (it may occur, for example, under a conformal inversion on a sphere). We call $\mathrm{dist}(f,x)$ *smooth* if it is a smooth function with respect to variations in f and x. We call it *symmetric*, if $\mathrm{dist}(f,x)$ is symmetric with respect to inversion of f, i.e., $\mathrm{dist}(f,x) = \mathrm{dist}(f^{-1}, f(x))$. A distortion is considered to be RDM if all three conditions of normalization, smoothness and symmetry are satisfied.

Normalized distortions follow a general expectation about the cost of "week" and "strong" space deformations, while symmetric distortions equally penalize on scaling and shrinking and thus are more stable during optimization. Minimizing non-symmetric measures may lead to undesirable shrinking and scaling if the boundary is not properly fixed. Consequently, non-symmetric measures are often substituted by their symmetric analogues, e.g., symmetric Dirichlet energy (14).

Using mathematical tools presented in Appendix A and Appendix B, the above distortion properties are generalized for entire space of deformation functions (see Definition C.1). Moreover, according to (C.2) and Corollary C.1, RDM measures approach infinity when an optimized vertex in Algorithm 2 crosses its 1-ring. Therefore, high values of the gradient pointing outward near ring boundaries avoid singularities, self intersections and flips of cells. Consequently, if we consider only RDM, then Algorithm 1 preserves orientation and local injectivity of f_0. This restriction has only limited impact on the algorithm applicability, since most non-RDM measures employed in the practice have "geometrically equivalent" RDM analogues. E.g., (16) and (13) can be substituted by (14) and (12), respectively.

Furthermore, according to (C.7) and (C.8) the gradient descent for RDMs is well defined even if a metric $\mathcal{D}$ has non-differentiable points and the process is stable for arbitrary constraints. Most importantly, according to [AL13], positively oriented simplicial maps are globally injective if they are injective on the boundary. Hence, in the RDM case, a sufficient condition for global injectivity are the positivity of an initial deformation and injectivity on the boundary $\partial\mathcal{S}$. If f_0 is positive and 1:1, then the second condition is obviously fulfilled for fixed boundary constraints. For more general constraints we employ a basic voxel traversing algorithm to find intersections between boundary faces during the line search stage of Algorithm 2. The length of a line search segment for unfixed boundary vertices is reduced if it intersects $\partial\mathcal{S}$ (see Fig. 7). Moreover, if $\mathcal{V}_i = \partial\mathcal{V}$ is incorporated into the mesh partition, then the above procedure can be integrated efficiently into parallel Algorithm 3.

On the other hand, the above property of RDMs prevents the gradient descent from repairing self intersections between cells if the initial map is not locally injective. In these cases a combination of the given distortion and other non-RDM penalty measures can be employed to repair self intersections and flips. Another approach to handling bad initializations is to perform alternating optimization of RDM and non-RDM measures.

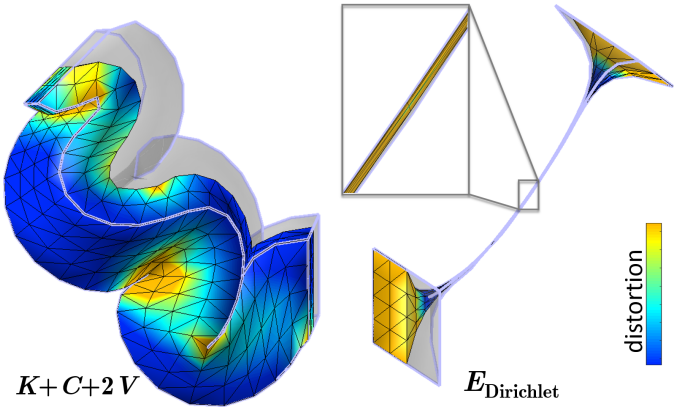

Fig. 7: Preserving global injectivity in tet meshes by employing ray tracing between boundary vertices.

Algorithm 3: ParallelGradDescent $(\mathcal{V}, \mathcal{V}_{\mathrm{fix}}, \mathcal{C}, f_0)$

Input:

- Source mesh $(\mathcal{V},\mathcal{C})$ and initial map f_0 fixed on $\mathcal{V}_{\mathrm{fix}}$.
- Mesh partition $\{(\mathcal{V}_i,\mathcal{C}_i)\}_{i=1}^k$ with initial maps $\{f_{0,i}\}_{i=1}^k$.

1. $\mathcal{V}_{\mathrm{com}} \leftarrow$ vertices that belong to at least two sub-meshes.
2. $\mathcal{C}_\partial \leftarrow$ cells that share $\mathcal{V}_{\mathrm{com}}$ vertices.
3. $\mathcal{V}_\partial \leftarrow$ vertices of $\mathcal{C}_\partial$.
4. $f_{0,\partial} \leftarrow f_0$ restricted to mesh $(\mathcal{V}_\partial,\mathcal{C}_\partial)$.
5. $j \leftarrow 0$.

repeat

$f_{j+1,\partial} \leftarrow \mathrm{Minimize}\big(f_{j,\partial}, \mathcal{V}_\partial, \mathcal{C}_\partial, \mathcal{V}_{\mathrm{fix}} \cup (\mathcal{V}_\partial/\mathcal{V}_{\mathrm{com}})\big)$;

for *i=1,...,k in parallel* **do**

Set $f_{j,i}$ to be equal to $f_{j,\partial}$ on $\mathcal{V}_{\mathrm{com}}$;

$f_{j+1,i} \leftarrow \mathrm{Minimize}(f_{j,i}, \mathcal{V}_i, \mathcal{C}_i, (\mathcal{V}_{\mathrm{fix}} \cup \mathcal{V}_{\mathrm{com}}) \bigcap \mathcal{V}_i)$;

Update $f_{j+1,\partial}$ on $\mathcal{V}_\partial/\mathcal{V}_{\mathrm{com}}$.

$j \leftarrow j+1$

until *Convergence of* $\sum_{i=1}^{k} \mathrm{dist}(f_{j,i})$;

6. $f_j \leftarrow$ Mapping of $(\mathcal{V},\mathcal{C})$ obtained from $f_{j,1},\ldots,f_{j,k}$.

Output: Mapping f_j that minimizes a given distortion.

4. Applications

4.1. Deformation of triangle and tetrahedral meshes

We first experiment with deformation of planar meshes. Fig. 8 demonstrates optimal deformations of an elephant model by applying Algorithm 1 with respect to various distortion measures under fixed boundary constraints. Given the source mesh

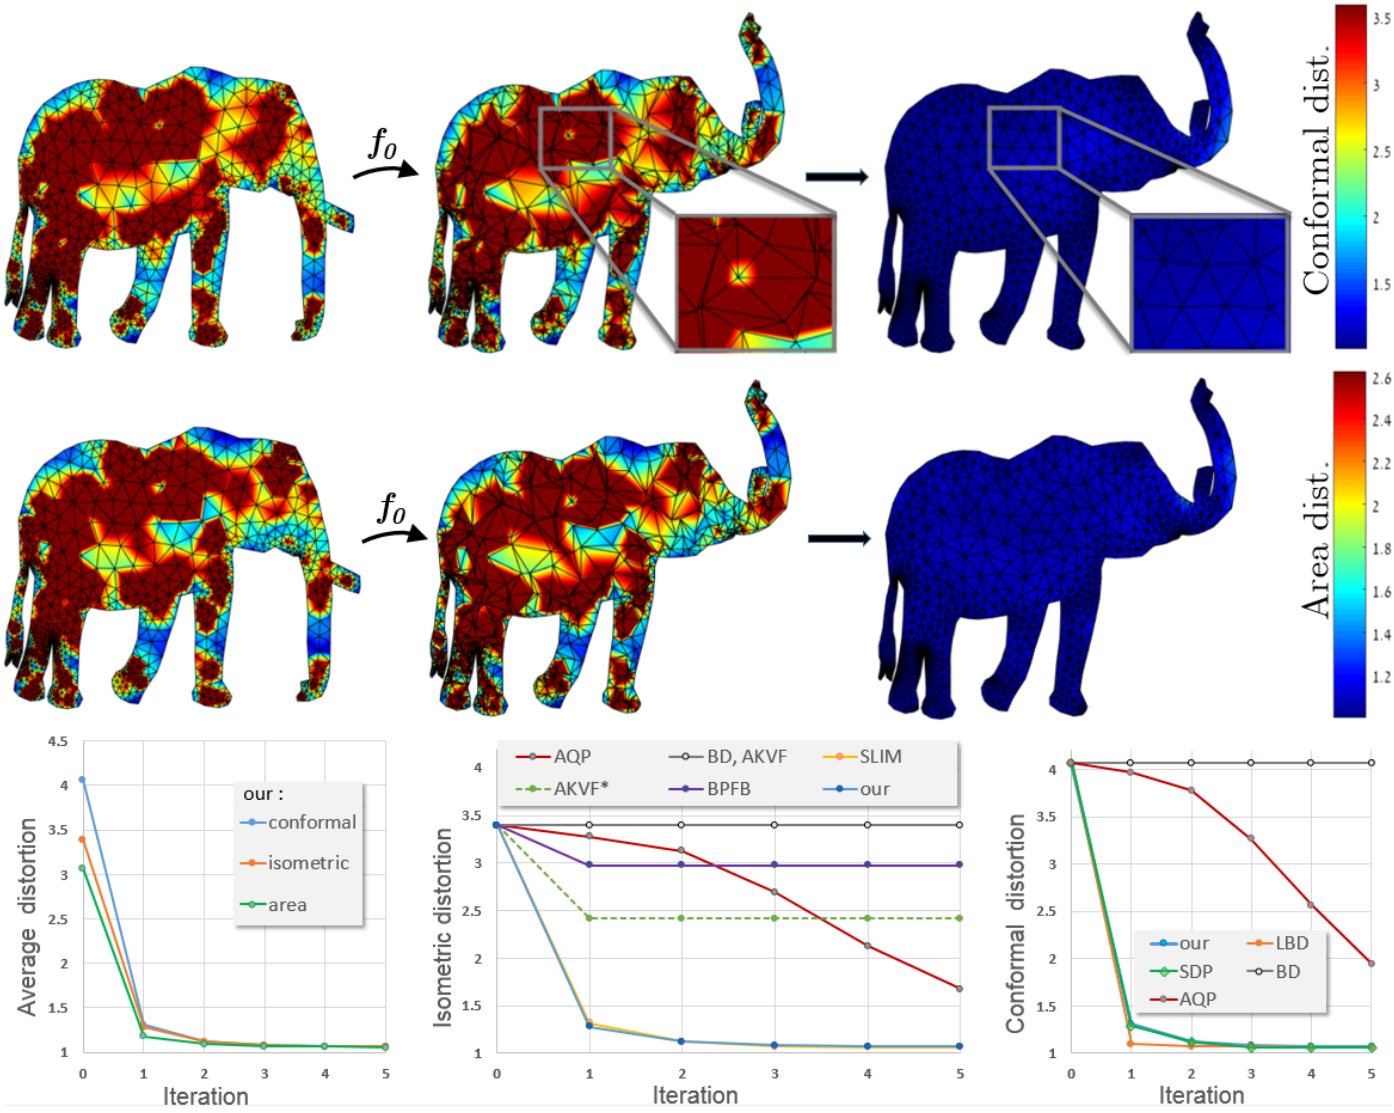

Fig. 8: Minimizing distortions for planar meshes. The boundary was fixed for all methods except AKVF* which denotes unconstrained AKVF (with free boundary). The source mesh, the initial map f_0 and the resulting deformations after 5 iterations of our Algorithm 1 are shown from left to right in the first two rows.

with m vertices, the initial deformation f_0 is computed as follows: using [CWKBC13] method we construct boundary-to-boundary mapping and then produce the target mesh with m vertices by triangulating the domain contained inside the boundary; f_0 is defined by a coarse correspondence between the source and target interior vertices. We also distort shapes of target triangles in $\mathrm{Im}(f_0)$ by adding a few iterations of a noise to positions of the interior vertices. As illustrated by Fig. 8 after few iterations of our algorithm we obtained global minimizers (up to negligible numeric errors) in all three cases, i.e., an area-preserving, isometric and conformal mappings, respectively. As emphasized in Fig. 8 the image of optimal deformation consists of almost regular triangles, meaning a high quality of the output mapping.

Note that fixed boundary constrains in the plane left no freedom for interior triangles to be resized during the optimization; hence output mappings for both conformal, and area ($V(f,x)$) distortions appear almost visually indistinguishable from each other. As shown in Fig. 8, our algorithm achieves better results than AQP for both isometric and conformal minimizations, similar results compared with LBD, SPD ans SLIM, whereas AKVF failed to reduce isometric distortion below its initial value under the fixed boundary and BD failed in both cases.

Next, we experiment with minimizing distortions on surfaces. As shown in Fig. 9, we use our algorithm to improve harmonic disc parametrization for non-regular meshes, where classical cotangent weighted harmonic maps produce artifacts. By randomly displacing target vertices of the face model we intensify initial distortions to extreme values (up to 10^{18}). Nevertheless, as depicted in Fig. 9, our algorithm attains almost perfect parameterization after 20 iterations. As shown in the logarithmically scaled distortion plot, other state-of-the-art methods fail or obtain significantly inferior results.

The main reason for poor results of AKVF, BPFB and SLIM in Figs. 8 and 9 is that these methods use global descent approaches. Although each of the methods employs a different algorithm to compute descend directions, all of them try to deform mesh elements simultaneously in a single direction.

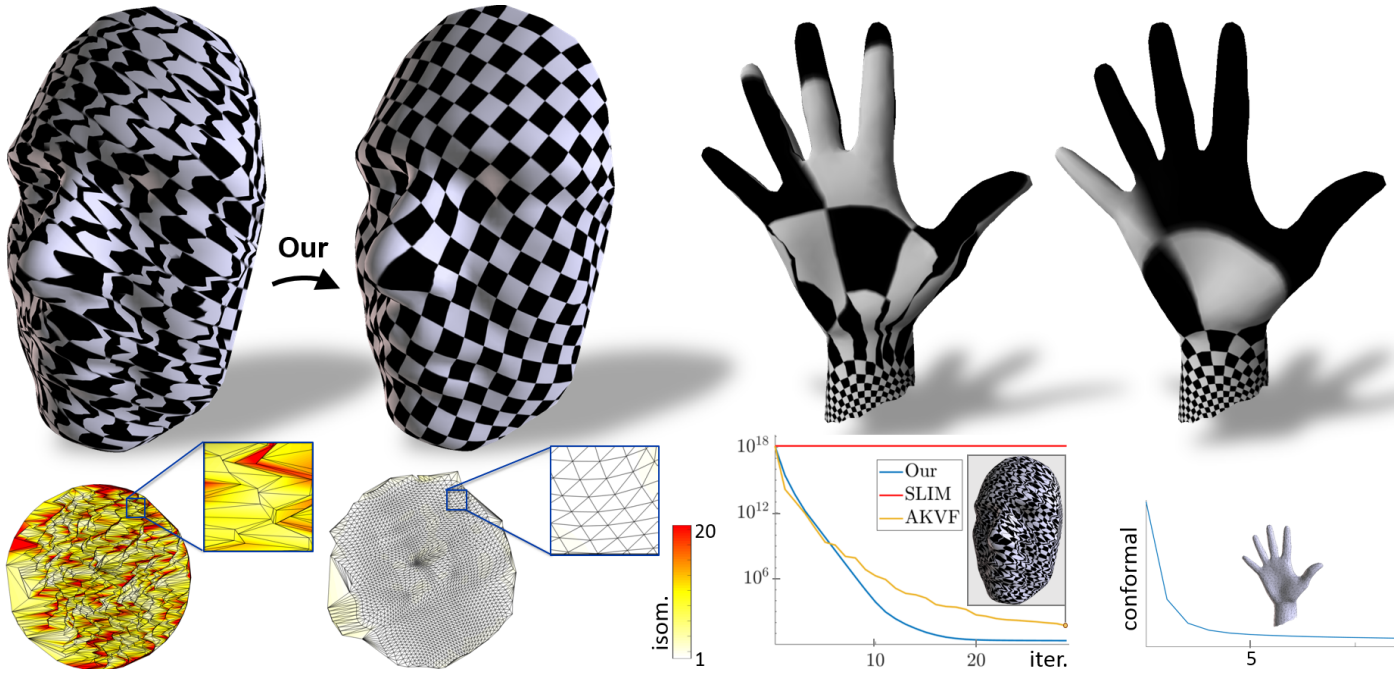

Fig. 9: Improving surface parametrizations via minimizing the product $K(f,x)C(f,x)$ and the conformal distortion $K(f,x)$ for the face and hand models, respectively. Multiplying isometric distortion by $K(f,x)$ results in a better alignment of vertices placed between small and large target triangles. The boundary is fixed to its initial position only for the hand model. The isometric distortion plot depicts $C(f,x)$ in the logarithmic scale.

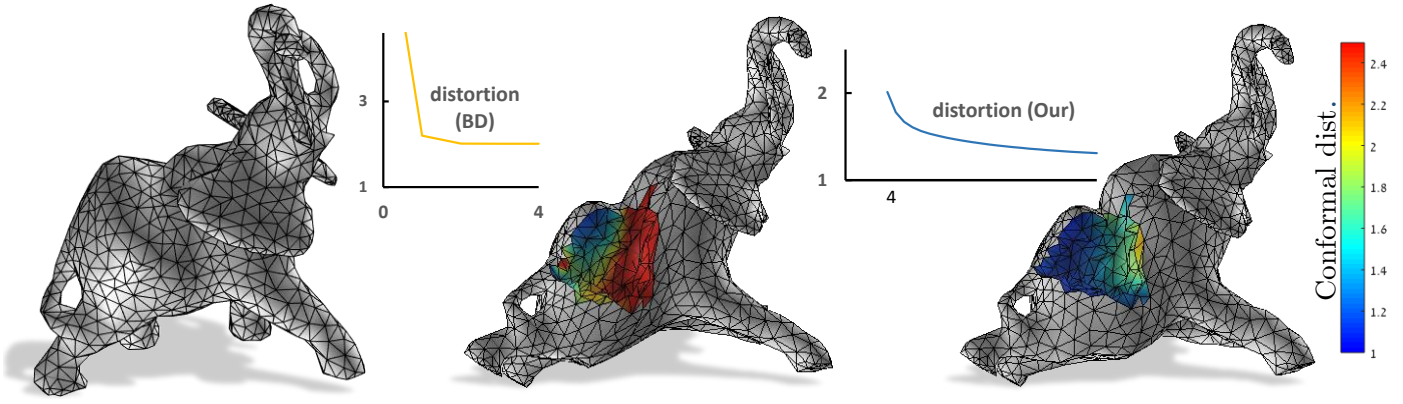

Fig. 10: Tet meshes of (from left to right) source domain, cross section of the image of the deformation obtained by BD method; cross section of the image of the deformation obtained by our algorithm, where the BD result is taken as the initial mapping. Optimizations were performed with fixed boundary constraints.

This strategy may work well when the optimization results in a smooth synchronized deformation. However, an noisy initialization produces large number of both regular and nearly singular simplices with chaotic gradients. In this case, a global descent process can be stuck easily in a line search stage. Since our algorithm is based on a local approach, it is much less sensitive to geometric noise and sparsity.

Finally, we test our method for deformations of volumetric models. As illustrated in Fig. 10, taking the result of [AL13] method as the initial deformation and then applying our algorithm significantly reduces conformal distortion. Fig. 5 depicts series of distortion minimizations for the initial radial mapping $f_0(x) = x\|x\|$ of a cube under fixed boundary constraints along with the plot of distortion values per each iteration. We compared our algorithm with state-of-the-art methods in Fig. 5 using the condition number (9) as a metric of conformal distortion, since except our work there is no method for minimizing more complicated metrics like $K(f,x)$ in 3D. As demonstrated by the figure, our method produced results that are favorable both visually and numerically. Furthermore, our technique can be particularly useful in modeling of surfaces and volumetric objects. Compared with related works, it applies to a wider set of geometric energies. As summarized in Figs. 12 and 7, employing different weights $\omega(v)$ and distortion measures along with their combinations produce a variety of volumetric shapes.

4.2. *Triangular and tetrahedral meshing*

By studding Figs. 5 and 8, our conclusion is that if a source mesh is regular and distortions $\mathrm{dist}(f_0, v)$ for boundary vertices are low, then minimizing distortions under fixed boundary constraints produces a nearly regular target mesh. To emphasize

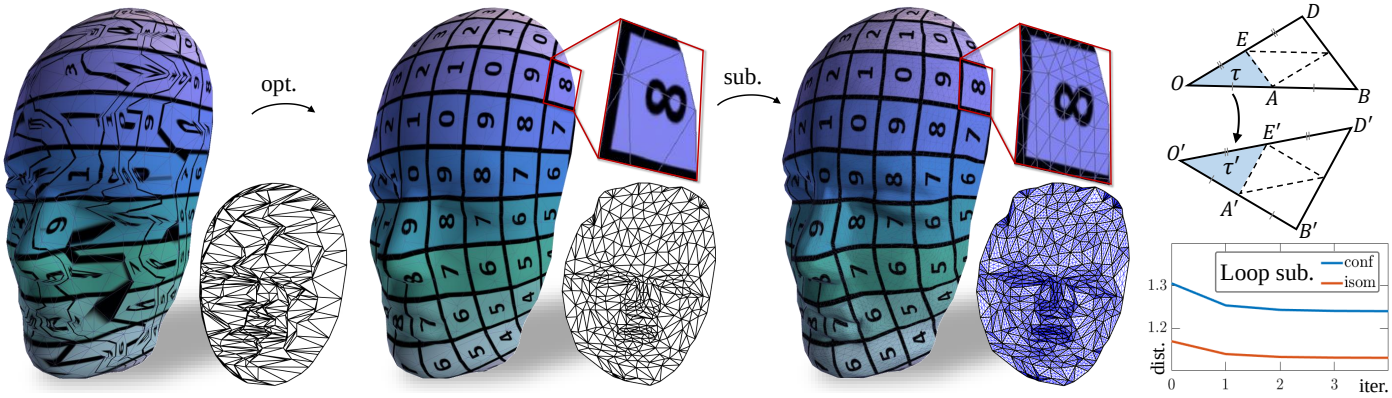

Fig. 11: Initial parametrization (left) is optimized (middle); next a high-resolution mapping obtained by repetitively subdividing the source and target triangles as depicted on top right. The plot (right bottom) shows distortions as functions of Loop subdivision iterations applied on the optimized map.

this observation, we compare resultant target meshes of a radial shape from Fig. 5 and a similar shape in 2D with state of the art methods of tetrahedral and triangular meshing, respectively. The results are shown in Fig. 15, were each pair of meshes contains nearly the same number of cells. We employed edge aspect ratio per vertex and the average tetrahedron aspect ratio as measures of the mesh quality. As shown by the figure our approach, based on ether isometric or conformal distortion minimizations, achieves favorable results.

4.3. Optimization and mesh subdivision

We tested the effect of surface subdivision on distortion optimization and found that subdividing both source and target meshes of an optimal map is a fast and accurate approximation of the optimal map in higher resolution. Indeed, most subdivision schemes tend to smooth out surface triangulation, leading to more regular triangles. A mapping between nearly regular simplices produces low conformal distortions due to shape similarities. If a subdivision process is scale invariant, then it preserves scale ratio between source and target triangles. Therefore, most subdivision methods preserve low conformal and low isometric distortions. In the basic linear subdivision scheme, shown in Fig. 11, each triangle is divided through its edge middle points. In this case, the resultant total distortion for any measure $\mathrm{dist}(f,x)$ is identical to the original one (for proof see Appendix D). Fig. 11 depicts an optimized face parametrization subdivided by this scheme. We also applied the Loop subdivision scheme [Loo87] on this example and found that it slightly reduces distortions (see the plot in Fig. 11). However, using a mesh subdivision directly for the optimization is impractical due to the exponential growth in the number of triangles.

4.4. Volumetric textures for surface mapping

Our method for optimizing tetrahedral domains can be employed on surfaces via a deformation of the enclosing volumes. This approach is highly advantageous on surfaces with complex topology, such as gyroid surfaces shown in Fig. 13. Gyroid-based structures have a great importance in 3D printing and material engineering applications (e.g., 3D graphene assembly [QJKB17]), but due to the high genus these surfaces cannot be mapped by standard techniques for triangle meshes. Even advance algorithms for n-genus manifolds are not efficient in this example.

However, consider the following procedure: construct a tet mesh $(\mathcal{V},\mathcal{C})$ that contains the triangulated surface S and represent surface vertices by their barycentric coordinates (18) with respect to their parent tetrahedrons. In other words, a surface is represented as a volumetric texture. Then, a simplicial map of $(\mathcal{V},\mathcal{C})$ constructed by our algorithm, can be effectively applied for S, resulting in a low distortion deformation of the surface into a desirable shape. Moreover, the runtime and accuracy of this procedure depend entirely on resolutions of S and $(\mathcal{V},\mathcal{C})$. Thus, its performance is independent of the geometric complexity. The first application illustrated in Fig. 13 is the mapping of cubic gyroid into a spherical shape induced by minimizing isometric distortion of the bounding volume. We use method [NSZ16] of star-shaped domain parametrization to construct the initial map. The second application illustrates a similar process of mapping volumetric structure from a cube to a head model. We employ in this process an inverse mass transportation [GLSY13] as the initial deformation.

Furthermore, as illustrated in Figs. 13 and 14 our algorithm can be used for improving existing surface and volumetric parametrization techniques, including basic methods that do not require heavy computations.

4.5. Implementation and performance

We have written our algorithm interface in MATLAB, while critical parts of the code were implemented in C++.

Table 1 summarizes the performance of our algorithms in comparison with related approaches. The behavior of related geometry optimization methods indicates that there exists a trade-off between geometric precision and the runtime. For example, according to Fig. 5 the closest result to our algorithm was achieved by [KABL14] method that suffers from a slow performance of SDP solvers.

Since significant part of our C++ code were automatically translated from MATLAB, we expect that a manual implementation of the entire algorithm in C/C++ will lead to much faster performance.

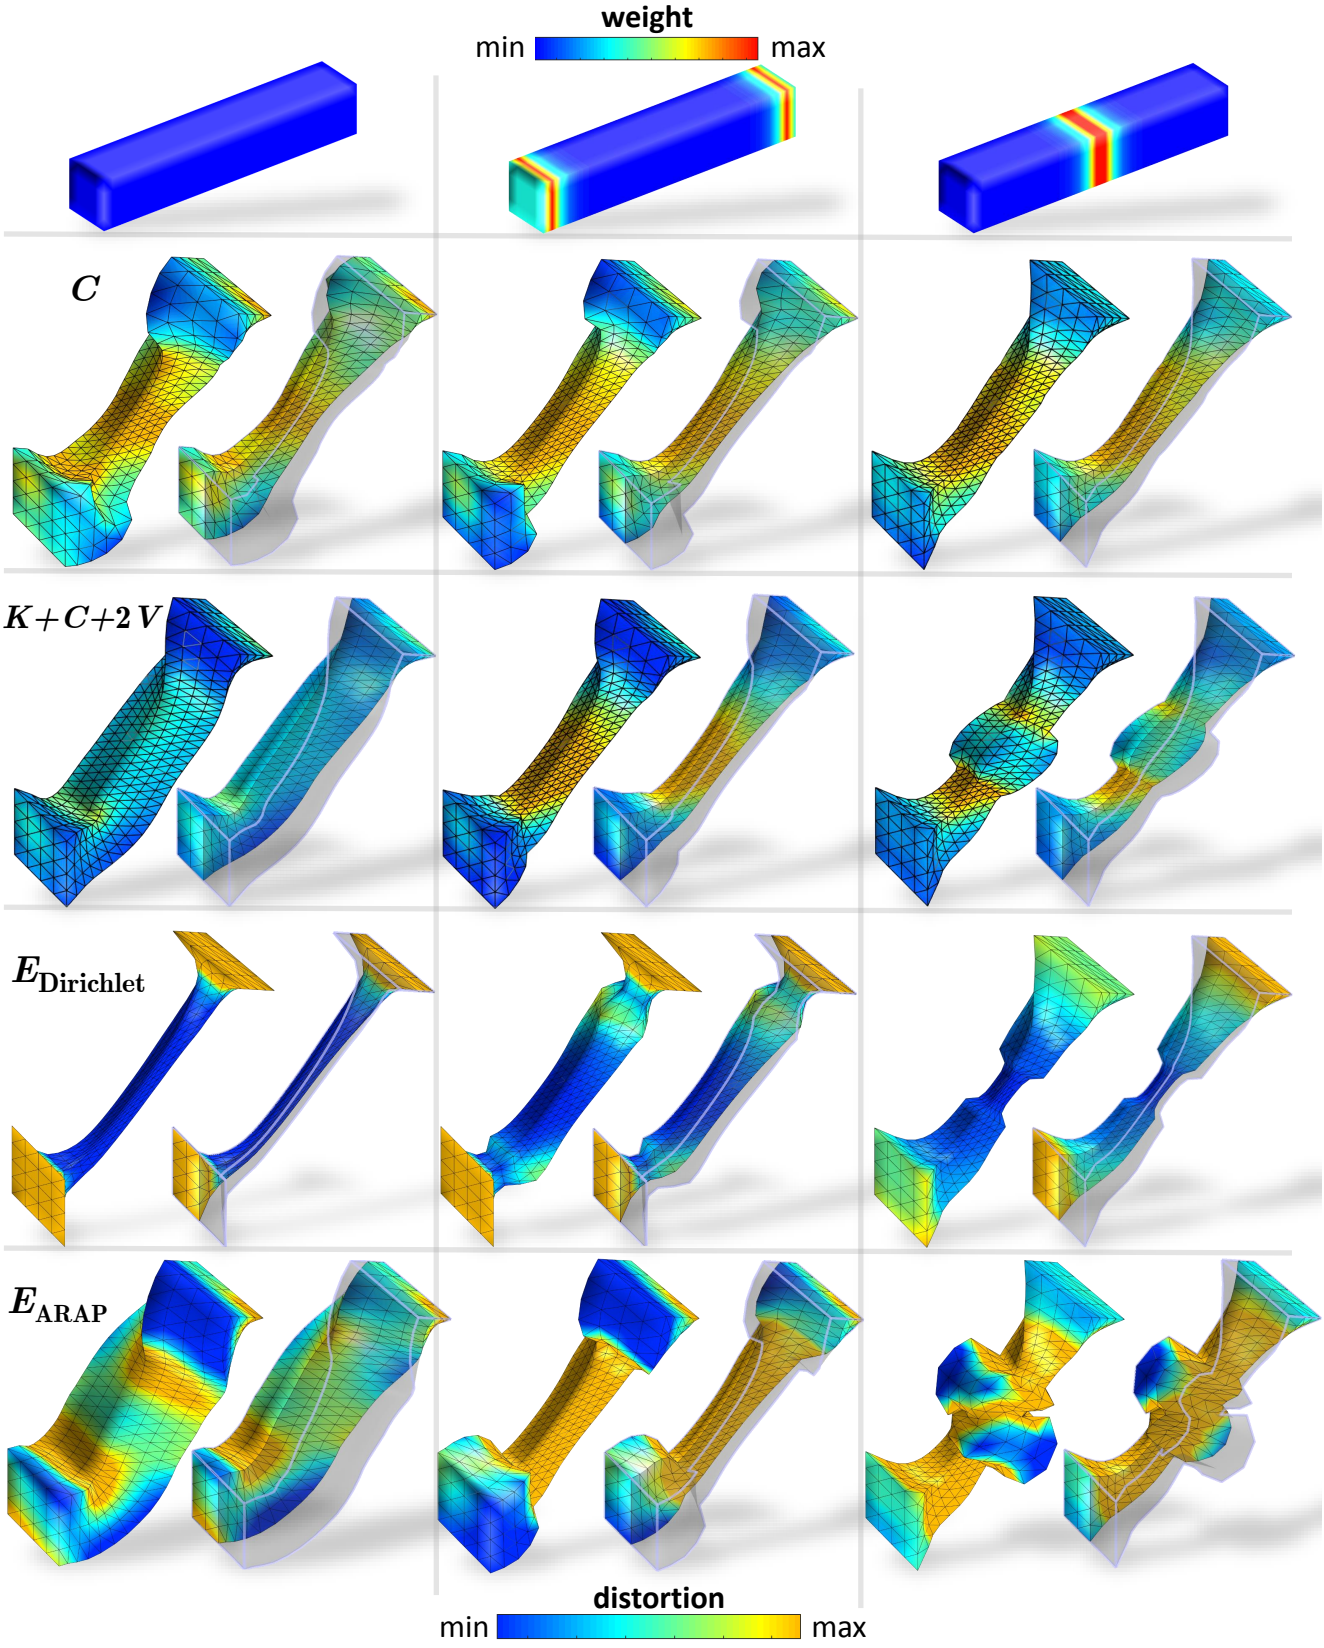

Fig. 12: Deformations of the volumetric bar (top) obtained via minimization of various distortion measures (shown per row) under constrained positions of bar's left and right edges. Columns correspond to different vertex weights: uniform, weights concentrated near the edges and weights concentrated at the middle.

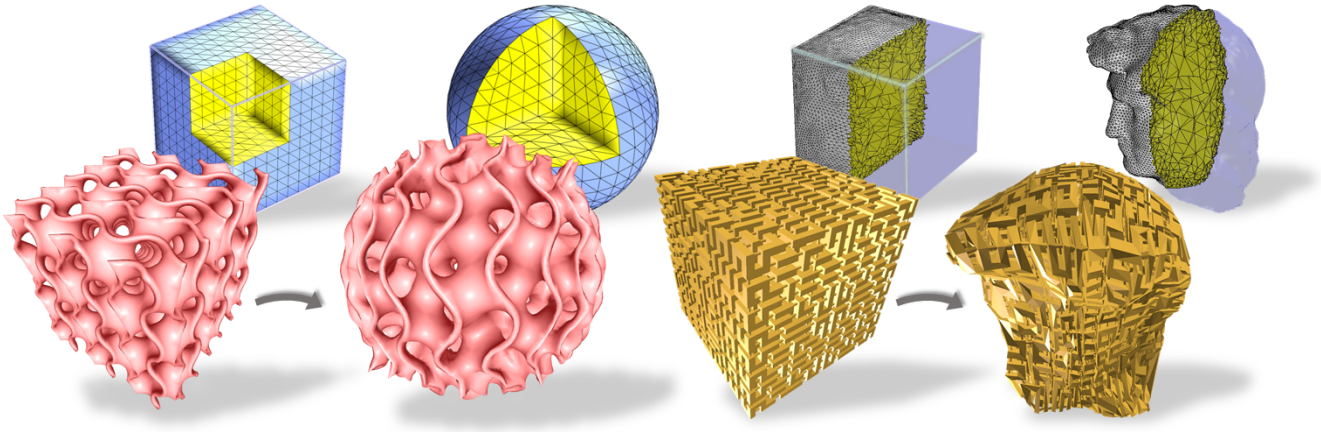

Fig. 13: Mapping of topologically complex gyroid surface (left) and generating a volumetric structure (right), using optimal deformations of the bounding volumes. Volumetric deformation for gyroid surface was obtained by minimizing isometric distortion under fixed "boundary-to-boundary" constraints for a parametrization of a bounding cube to a ball (left top). Volumetric deformation of a cube to the head model was obtained via minimization of conformal distortion for inverse mass transportation of tetrahedral mesh to a cube (right top).

5. Conclusion

We have presented a framework for mapping and deforming discrete geometric data based on minimization of distortion measures. Our algorithm is easy to reproduce for a parallel run, and its implementation does not require external optimization tools like SDP and large scale linear solvers employed in related methods. We have shown that our technique can outperform existing algorithms in terms of runtime performance and geometric accuracy. Moreover, our algorithm can be further improved in a straightforward manner by employing analytic gradient approximations and by optimizing parallel computations as discussed in Section 3.1.

Despite the relative simplicity of our approach, the underlying steepest descent techniques were not analyzed properly in the previous studies presented in the literature. Our work closes the gap trough a formal analysis of deformations and fundamental geometric measures involved in the optimization. Our approach is very general and can be extended to higher dimensions and by employing local coordinates to manifolds.

We employed our method so far in a collection of geometry processing problems. We expect to find more applications in related fields, including medical imaging and computer vision. In particular, we expect that computing distortions of optimal deformations between 3D objects can be applied in change detections and geometric similarity assessments.

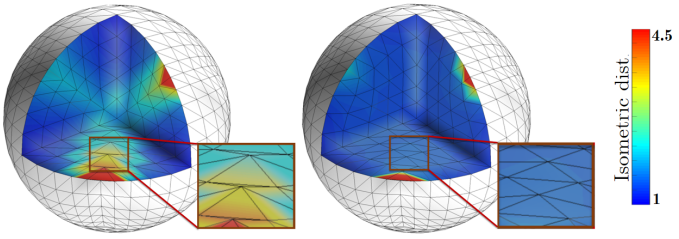

Fig. 14: Parametrization to a ball from Fig. 13: method of [NSZ16] (left) and parametrization improved by Algorithm 1 (right).

Acknowledgment

Research has been supported in part by the office of the Israeli Chief Scientist–OMEK Consortium and by the Ollendorff Minerva Center for Vision and Image Sciences.

Table 1: Conformal optimization for meshes of 0.8K (shown in Fig.5), 6K and 50K tets. Runtime measures for Intel i5-5590 (4 Cores) are shown in seconds for Algorithm 1 (1 thread) and for Algorithm 3 for parallel processing of 2 and 4 segments. Note, that comparison with BD is unfair, since its numerical results are much worse than ours. Our algorithm achieves the same results for 6K mesh after a one second. We did not compare runtime with AQP and LBD, since these methods failed to reduce distortions below the initial level.

Method	Tet mesh	1 thread	2 threads	4 threads
Our	0.8K	0.2	0.2	0.7
	6K	11.2	6.5	7.4
	50K	107	59.6	49.4
BD	0.8K	4.8		
	6K	11.5		
	50K	120		
SDP	0.8K	33.7		
	6K	328		
	50K	2235		

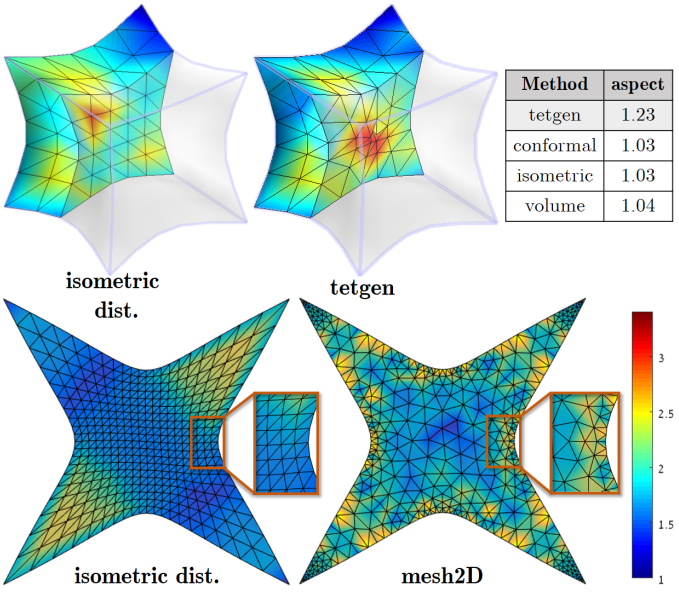

Fig. 15: Meshes produced via minimizing distortions by Algorithm 1 (left) and state-of-the-art algorithms "tetgen" [Si09] and "mesh2D" [Eng09] applied for obtaining tet and triangle meshes, respectively. The table contains average cell aspect ratios and colors are according to edge aspects.

Appendix A. Formal definitions

Deformation functions and smooth deformations considered in Section 3 are formally defined as follows:

Definition A.1. Let $\mathcal{S}$ be a *proper domain* of $\mathbb{R}^n$, namely, $\mathcal{S} \subset \mathbb{R}^n$ is a compact set with a non-empty interior. Then, a continuous function $f : \mathcal{S} \to \mathbb{R}^n$ is called a *local diffeomorphism* (or *smooth deformation*) if for each $x \in \mathrm{int}(\mathcal{S})$ there is a neighborhood, called a *local neighborhood of f at x*, where f is a smooth bijective mapping with a smooth inverse.

Similarly, a continuous function f is called a *local homeomorphism* (or a *deformation*) if for each $x \in \mathrm{int}(\mathcal{S})$ there is a *local neighborhood* N_x such that $f|_{N_x}$ is a homeomorphism, i.e., $f|_{N_x}$ is continuous bijective mapping.

We denote by $\mathrm{Hom}(\mathbb{R}^n)$ and $\mathrm{Def}(\mathbb{R}^n)$ the set of deformation functions and the set of smooth deformation functions of proper domains in $\mathbb{R}^n$, respectively. The interior of a proper domain of a function $f \in \mathrm{Hom}(\mathbb{R}^n)$ will be denoted by $\mathrm{Dom}(f)$.

Following the above notations, distortion measures are formally defined as functions

$$\mathrm{dist} : \{ (f,x) | f \in \mathrm{Hom}(\mathbb{R}^n),\ x \in \mathrm{Dom}(f) \} \to \mathbb{R},$$

and the *first order precision* property of Definition 2.1 is formalized as follows: Assume that $f, g \in \mathrm{Hom}(\mathbb{R}^n)$ coincide on x_0 and have a common local neighborhood N of x_0. If $f(x) - g(x) = o(\|x - x_0\|)$ in N, then $\mathrm{dist}(f,x_0) = \mathrm{dist}(g,x_0)$.

Appendix B. Canonical representations of distortions

Lemma 1. *The Jacobian of $f \in \mathrm{Def}(\mathbb{R}^n)$ is a full rank matrix, and it therefore contains only positive singular values, denoted in the descending order by $\sigma_1(df_x), \ldots, \sigma_n(df_x)$, where df_x is the Jacobian matrix of f at x.*

Based on above definitions and Lemma 1, Theorem 2.1 can be formally stated and proven as follows:

Theorem B.1. *For $f \in \mathrm{Def}(\mathbb{R}^n)$, $x \in \mathrm{Dom}(f)$, a first order distortion measure, $\mathrm{dist}(f,x)$, can be expressed as a function of singular values of the Jacobian df_x, i.e.,*

$$\exists \mathcal{D} : \mathbb{U}^n \to \mathbb{R}, \mathrm{dist}(f,x) = \mathcal{D}(\sigma_1(df_x), \ldots, \sigma_n(df_x)),$$

where $\mathbb{U}^n$ is defined by (5).

Proof. First, let us show that $\mathrm{dist}(f,x_0) = \mathcal{D}(df_{x_0})$, that is, $\mathrm{dist}(f,x_0)$ is a function of the entries of the Jacobian matrix at x_0. Since distortion measures are invariant to compositions with rigid transformations, we assume w.l.o.g. that $x_0 = f(x_0) = 0$. Let N_0 be a local neighborhood of f at 0 and consider the linear function $g(x) = df_0 x$ for $x \in N_0$. Next, since f is diffeomorphism in N_0 it can be linearly approximated by the first term of its Taylor series; hence $f(x) - g(x) = o(\|x\|)$. Therefore, according to the 1st order precision property

$$\mathrm{dist}(f,0) = \mathrm{dist}(df_0(\cdot),0),$$

where $df_0(\cdot)$ denotes the linear function $x \mapsto df_0 \cdot x$. Consequently, $\mathrm{dist}(f,0)$ is a function of the entries of Jacobian df_0.

Finally, let $df_0 = U\Sigma V^T$ be SVD of the Jacobian. Then, the second property of Definition 2.1 applied for rigid transformations $T_1 = V, T_2 = U^T$ yields

$$\mathrm{dist}(df_0(\cdot),0) = \mathrm{dist}(\Sigma(\cdot),0),$$

where Σ is a diagonal matrix $\mathrm{diag}(\sigma_1(df_0), \ldots, \sigma_n(df_0))$ and the singular values $(\sigma_1(df_0), \ldots, \sigma_n(df_0))$, put in the descending order, constitute, according to Lemma 1, an element of $\mathbb{U}^n$. □

Appendix C. Regular distortion measures

Definition C.1. Let $\mathrm{dist}(\cdot,\cdot)$ be distortion measure and let $\mathcal{D}$ be its canonical representation. We call $\mathrm{dist}(\cdot,\cdot)$ a *a regular distortion measure (RDM)* if the following properties are satisfied:

1. **Normalization.** Distortion measure $\mathrm{dist}(f,x)$ is called *normalized* if $\mathrm{dist}(f,x) \in [m,\infty)$ for any $f \in \mathrm{Def}(\mathbb{R}^n)$, $x \in \mathrm{Dom}(f)$ and the following conditions are met:

$$\mathrm{dist}(Id_D,x) \equiv m\,, \forall x \in \mathrm{int}(D)\,, \tag{C.1}$$

$$\mathcal{D}(\sigma_1,\ldots,\sigma_n) \to \infty \text{ when } \sigma_n \to 0 \text{ and } \sigma_1 > \varepsilon > 0; \tag{C.2}$$

   where Id_D denotes the identity function of a proper domain D. Usually, minimal distortion values are $m = 0$ or $m = 1$.

2. **Symmetry**. We call $f^\dagger$ *a local inversion* of deformation f, if for each $x \in \mathrm{Dom}(f)$ there are local neighborhoods X of f at x, and Y of $f^\dagger$ at $y = f(x)$, s.t., $(f|_X)^{-1} = f^\dagger|_Y$. If for any $f \in \mathrm{Def}(\mathbb{R}^n)$

$$\mathrm{dist}(f,x) = \mathrm{dist}\left(f^\dagger, f(x)\right), \forall x \in \mathrm{Dom}(f)\,, \tag{C.3}$$

   then $\mathrm{dist}(\cdot,\cdot)$ is called *symmetric*. In other words, symmetric distortions do not distinguish between mapping a source into a target domain and the vice versa deformation.

3. **Smoothness**: A distortion measure $\mathrm{dist}(f,x)$ is called smooth if $\mathcal{D}$ is a smooth function a.e. (almost everywhere).

Corollary C.1. *Immediate consequences of Definition C.1 are:*

1. *Each of the following properties is equivalent to (C.1):*

$$\mathcal{D}(1,1,...,1) = m\,, \tag{C.4}$$

$$\forall T\text{-}rigid : \mathrm{dist}(T,x) \equiv m\,. \tag{C.5}$$

2. *The symmetry is equivalent to*

$$\mathcal{D}(\sigma_1,\ldots,\sigma_n) = \mathcal{D}\left(\frac{1}{\sigma_n},\ldots,\frac{1}{\sigma_1}\right) \forall \sigma \in \mathbb{U}^n\,. \tag{C.6}$$

3. *If* $\mathrm{dist}(\cdot,\cdot)$ *is symmetric, then (C.2) implies*

$$\mathcal{D}(\sigma_1,\ldots,\sigma_n) \to \infty \text{ when } \sigma_1 \to \infty \text{ and } \sigma_n < N < \infty. \tag{C.7}$$

Proof. By Definition 2.1 and Theorem B.1, $\mathrm{dist}(T,x) = \mathrm{dist}(T^{-1}\circ T,x) = \mathrm{dist}(Id_D,x) = \mathcal{D}(1,...,1)$. Therefore, (C.1), (C.4) and (C.5) are equivalent. The second proposition follows from the fact that if $\sigma_1,\ldots,\sigma_n$ are singular values of full rank df_x, then $\sigma_n^{-1},\ldots,\sigma_1^{-1}$ are singular values of df_x^{-1}. The third proposition is the immediate consequence of (C.6). □

Intuitively, we expect a small discrepancy in $\mathrm{dist}(f,x)$ when ether f or x are slightly changed. Clearly, if $\mathrm{dist}(\cdot,\cdot)$ is smooth, then it is continuous in both f and x [3]. We require C^1 continuity of $\mathcal{D}$ to ensure that the gradient descent processes is well defined. Since many known distortion measures are set as a maximum of several functions, it is important to relax the smoothness condition to be defined almost everywhere. In this case, we can define the gradient even for a non-smooth point x using the *weak derivative*

$$\left.\frac{\partial\,\mathrm{dist}(f,\cdot)}{\partial h}\right|_x \triangleq \limsup_{\varepsilon\to 0}\frac{\mathrm{dist}(f,x+\varepsilon h)-\mathrm{dist}(f,x)}{\varepsilon\|h\|}\,. \tag{C.8}$$

Appendix D. Subdivision and distortion preservation

Denote by f and g the original map and the map subdivided by the scheme shown in Fig. 11 (top-right), respectively.

Proposition D.1. *$E(f) = E(g)$, where E is th energy (25) with the weights $w(c) = \mathrm{area}(c)$.*

Proof. Following the annotation of Fig. 11, denote by c, c' and τ, τ' two pairs of source-target cells that belong to original and subdivided meshes, respectively. The triangle middle point theorem implies that c and c' are divided into congruent triangles, respectively. Therefore, $\mathrm{dist}(g_\tau)$ are identical for any τ contained in c and it is enough to proof $\mathrm{dist}(g_\tau) = \mathrm{dist}(f_c)$ for τ chosen as in Fig. 11. Assume w.l.o.g. $O = O' = (0,0,0)$, then g_τ and f_c are linear and $g_\tau(B) = g_\tau(2A) = 2g_\tau(A) = B'$. Similarly, $g_\tau(D) = D'$ which implies $g_\tau = f_c$ and $w(c)\,\mathrm{dist}(f_c) = \sum_{\tau\subset c} w(\tau)\,\mathrm{dist}(g_\tau)$ for any cell c, thus $E(f) = E(g)$. □

References

[Ahl66] AHLFORS L. V.: *Complex Analysis an Introduction to the Theory of Analytic Functions of One Complex Variable*. 1966.

[AL13] AIGERMAN N., LIPMAN Y.: Injective and bounded distortion mappings in 3D. *ACM Trans. Graph 32*, 4 (2013), 106:1–106:14.

[APL14] AIGERMAN N., PORANNE R., LIPMAN Y.: Lifted bijections for low distortion surface mappings. *ACM Transactions on Graphics (TOG) 33*, 4 (2014), 69.

[BCWG09] BEN-CHEN M., WEBER O., GOTSMAN C.: Variational harmonic maps for space deformation. In *ACM Transactions on Graphics (TOG)* (2009), vol. 28, ACM, p. 34.

[BDS*12] BOUAZIZ S., DEUSS M., SCHWARTZBURG Y., WEISE T., PAULY M.: Shape-up: Shaping discrete geometry with projections. In *Computer Graphics Forum* (2012), vol. 31, Wiley Online Library, pp. 1657–1667.

[Car74] CARAMAN P.: *n-dimensional quasiconformal (QCF) mappings. Revised, enlarged and translated from the Roumanian by the author.* 1974.

[CBSS17] CLAICI S., BESSMELTSEV M., SCHAEFER S., SOLOMON J.: Isometry-aware preconditioning for mesh parameterization. In *Computer Graphics Forum* (2017), vol. 36, Wiley Online Library, pp. 37–47.

[CPS15] CHERN A., PINKALL U., SCHRÖDER P.: Close-to-conformal deformations of volumes. *ACM Transactions on Graphics (TOG) 34*, 4 (2015), 56.

[CWKBC13] CHEN R., WEBER O., KEREN D., BEN-CHEN M.: Planar shape interpolation with bounded distortion. *ACM Transactions on Graphics (TOG) 32*, 4 (2013), 108.

[3] $\sigma_i(df_x)$ are continuous in x and f if f is restricted to $\mathrm{Def}(\mathbb{R}^n)$.

[Eng09] ENGWIRDA D.: Mesh2d-automatic mesh generation. *URL: http://www. mathworks. com/matlabcentral/fileexchange/25555-mesh2d-automatic-mesh-generation* (2009).

[FLG15] FU X.-M., LIU Y., GUO B.: Computing locally injective mappings by advanced mips. *ACM Transactions on Graphics (TOG) 34*, 4 (2015), 71.

[GLSY13] GU X., LUO F., SUN J., YAU S.-T.: Variational principles for minkowski type problems, discrete optimal transport, and discrete monge-ampere equations. *arXiv preprint arXiv:1302.5472* (2013).

[Hor01] HORMANN K.: *Theory and applications of parameterizing triangulations*. PhD thesis, Citeseer, 2001.

[KABL14] KOVALSKY S. Z., AIGERMAN N., BASRI R., LIPMAN Y.: Controlling singular values with semidefinite programming. *ACM Transactions on Graphics (TOG) 33*, 4 (2014), 68.

[KABL15] KOVALSKY S. Z., AIGERMAN N., BASRI R., LIPMAN Y.: Large-scale bounded distortion mappings. *ACM Trans. Graph 34*, 6 (2015), 191.

[KGL16] KOVALSKY S. Z., GALUN M., LIPMAN Y.: Accelerated quadratic proxy for geometric optimization. *ACM Transactions on Graphics (TOG) 35*, 4 (2016), 134.

[LGW*07] LI X., GUO X., WANG H., HE Y., GU X., QIN H.: Harmonic volumetric mapping for solid modeling applications. In *Proceedings of the ACM* (2007), pp. 109–120.

[LLL15] LEE Y. T., LAM K. C., LUI L. M.: Landmark-matching transformation with large deformation via n-dimensional quasi-conformal maps. *J. Sci. Comput.* (2015), 1–29.

[Loo87] LOOP C.: Smooth subdivision surfaces based on triangles.

[LPRM02] LÉVY B., PETITJEAN S., RAY N., MAILLOT J.: Least squares conformal maps for automatic texture atlas generation. In *Acm transactions on graphics (tog)* (2002), vol. 21, ACM, pp. 362–371.

[NCQ*18] NAITSAT A., CHENG S., QU X., FAN X., SAUCAN E., Y. Z. Y.: Geometric approach to detecting volumetric changes in medical images. *Journal of Computational and Applied Mathematics 329* (2018), 37–50.

[NSY17] NAITSAT A., SAUCAN E., Y. Z. Y.: A differential geometry approach for change detection in medical imaging. In *Computer-Based Medical Systems. CBMS'17* (2017), IEEE.

[NSZ15] NAITSAT A., SAUCAN E., ZEEVI Y. Y.: Volumetric quasi-conformal mappings - quasi-conformal mappings for volume deformation with applications to geometric modeling. In *Proceedings of VISIGRAPP 2015* (2015), pp. 46–57.

[NSZ16] NAITSAT A., SAUCAN E., ZEEVI Y. Y.: Geometric approach to estimation of volumetric distortions. In *Proceedings of VISIGRAPP 2016* (2016).

[PP12] PAILLÉ G.-P., POULIN P.: As-conformal-as-possible discrete volumetric mapping. *Computers & Graphics 36*, 5 (2012), 427–433.

[QJKB17] QIN Z., JUNG G. S., KANG M. J., BUEHLER M. J.: The mechanics and design of a lightweight three-dimensional graphene assembly. *Science advances 3*, 1 (2017), e1601536.

[RPPSH17] RABINOVICH M., PORANNE R., PANOZZO D., SORKINE-HORNUNG O.: Scalable locally injective mappings. *ACM Transactions on Graphics (TOG) 36*, 2 (2017), 16.

[SA07] SORKINE O., ALEXA M.: As-rigid-as-possible surface modeling. In *Symposium on Geometry processing* (2007), vol. 4.

[SBCBG11] SOLOMON J., BEN-CHEN M., BUTSCHER A., GUIBAS L.: As-killing-as-possible vector fields for planar deformation. In *Computer Graphics Forum* (2011), vol. 30, Wiley Online Library, pp. 1543–1552.

[SCL*17] SU K., CHEN W., LEI N., ZHANG J., QIAN K., GU X.: Volume preserving mesh parameterization based on optimal mass transportation. *Computer-Aided Design 82* (2017), 42–56.

[SCQ*16] SU K., CUI L., QIAN K., LEI N., ZHANG J., ZHANG M., GU X. D.: Area-preserving mesh parameterization for poly-annulus surfaces based on optimal mass transportation. *Computer Aided Geometric Design 46* (2016), 76–91.

[SDGT12] SANDHU R., DOMINITZ A., GAO Y., TANNENBAUM A.: Volumetric mapping of genus zero objects via mass preservation. *arXiv preprint arXiv:1205.1225* (2012).

[Si09] SI H.: *A Quality Tetrahedral Mesh Generator and Three-Dimensional Delaunay Triangulator*. Tetgen, 2009. URL: http://tetgen.berlios.de.

[SS15] SMITH J., SCHAEFER S.: Bijective parameterization with free boundaries. *ACM Transactions on Graphics (TOG) 34*, 4 (2015), 70.

[V71] VÄISÄLÄ J.: *Lectures on n-Dimensional Quasiconformal Mappings*. Springer-Verlag Berlin Heidelberg New York, 1971.

[WGY*03] WANG Y., GU X., YAU S.-T., ET AL.: Volumetric harmonic map. In *Communications in Information & Systems* (2003), vol. 3, pp. 191–202.

[WN99] WRIGHT S. J., NOCEDAL J.: Numerical optimization. *Springer Science 35*, 67-68 (1999), 135–142.

[ZGZ*14] ZHANG M., GUO R., ZENG W., LUO F., YAU S.-T., GU X.: The unified discrete surface ricci flow. *Graphical Models 76*, 5 (2014), 321–339.

[ZLYG09] ZENG W., LUO F., YAU S.-T., GU X. D.: Surface quasi-conformal mapping by solving beltrami equations. In *IMA International Conference on Mathematics of Surfaces* (2009), Springer, pp. 391–408.