



On Inversion-Free Mapping and Distortion Minimization

Alexander Naitzat¹ · Gregory Naitzat² · Yehoshua Y. Zeevi¹

Received: 13 June 2020 / Accepted: 11 May 2021 / Published online: 6 June 2021
© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2021

Abstract

This paper addresses a general problem of computing inversion-free maps between continuous and discrete domains that induce minimal geometric distortions. We will refer to this problem as *optimal mapping* problem. Finding a good solution to the optimal mapping problem is a key part in many applications in geometry processing and computer vision, including: parameterization of surfaces and volumetric domains, shape matching and shape analysis. The first goal of this paper is to provide a self-contained exposition of the optimal mapping problem and to highlight the interrelationship of various aspects of the problem. This includes a formal definition of the problem and of the related unitarily invariant geometric measures, which we call *distortions*. The second goal is to identify novel properties of distortion measures and to explain how these properties can be used in practice. Our major contributions are: (i) formalization and juxtaposition of key concepts of the optimal mapping problem, which so far have not been formalized in a unified manner; (ii) providing a detailed survey of existing methods for optimal mapping, including exposition of recent optimization algorithms and methods for finding injective mappings between meshes; (iii) providing novel theoretical findings on practical aspects of geometric distortions, including the multi-resolution invariance of geometric energies and the characterization of convex distortion measures. In particular, we introduce a new family of convex distortion measures, and prove that, on meshes, most of the existing distortion energies are non-convex functions of vertex coordinates.

Keywords Distortion · Non-linear optimization · Inversion-free mapping · Conformal mapping · Surface parameterization · Volumetric mapping

1 Introduction

One of the central issues in geometry processing and imaging is how to incur minimal distortion to original shape when deforming it to satisfy certain geometric constraints. This problem is often solved by optimization of geometric energies defined in a finite element fashion: the original shape is encoded by a simplicial complex, its global deformation is described by a simplicial map on that complex and the shape distortion is computed as a sum of local distortions over individual simplices. The goal is to find a simplicial map that minimizes the shape distortion and induces a consistent orientation of simplices (*inversion-free mapping*).

For example, consider a texture mapping problem — a process in which an RGB image is wrapped around a three dimensional object. This problem is also referred to as the surface parameterization or surface flattening. To illustrate, given a piece of cloth with an image printed on it, we want to stretch it and fold it around a triangulated surface in \mathbb{R}^3 causing as few distortions to the printed image as possible. The same problem can be generalized from surfaces to volumetric domains. For instance, consider parameterization of a three dimensional object, obtained by mapping its interior onto a simple domain, such as a cube or a ball (a *canonical domain*). In many fields, including machine learning and medical imaging, data, sampled on surfaces and volumetric regions, are parameterized and mapped onto a canonical domain. This way one can standardize geometric data across different samples.

For example, it is simpler to compare segments of CT scans by mapping these segments into a common canonical domain, where changes in geometric features can be detected more easily (see Fig. 1). Similarly, in machine learning, it is

✉ Alexander Naitzat
anaitzat@campus.technion.ac.il

¹ Viterbi Faculty of Electrical Engineering, Technion - Israel Institute of Technology, Haifa, Israel

² Department of Statistics, University of Chicago, Chicago, IL, USA

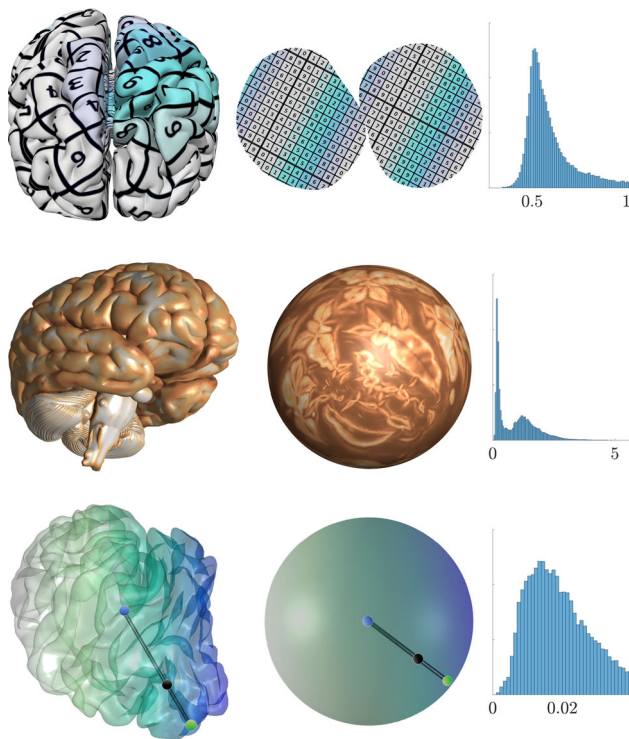


Fig. 1 Mapping complex 3D objects into the plane and onto canonical domains to get a simplified representation of these objects. At the top, we show texture mapping of the brain surface obtained by solving the distortion minimization problem on the triangle mesh. To get a continuous map, we cut the source domain into a disc-topology surface and minimize isometric distortion, induced by the texture mapping. In the middle, we show the mapping of genus-0 triangle mesh onto a sphere via the linear harmonic method [19]. At the bottom, we extend harmonic mapping of the surface of the right hemisphere into its volume, represented by a tetrahedral mesh. The volumetric map is computed by, first, stretching rays between the origin and boundary vertices according to their mutual distances (radial stretching method [78]), and then, by optimizing isometric distortion over interior tetrahedra. On the right side, we depict histograms of average singular values (78) (approximate conformal factors) of maps between simplices

simpler to train a neural network on regular data, such as a collection of RGB images, than training it on triangulated surfaces. By using texture mapping, one can map surfaces into images and therefore, one can generalize neural network architectures, designed for images, to surfaces (see Fig. 2).

Shape matching is another problem that involves mapping between triangulated domains with low geometric distortions. If a mapping preserves essential geometric features, then it can be used to transfer data between multiple domains. For example, shape matching methods can be used to transfer textures between surfaces, labels between volumetric medical images and etc.

To summarize, in all these examples, our task is to compute a “nice” deformation of a compact subset of Euclidean space that minimizes selected distortion criteria and whose image satisfies certain geometric constraints.

As can be anticipated, “nice” is not a universal property but a task related notion, such as visual image distortion avoidance, map injectivity and etc. In practice, very many physically motivated distortion criteria can be formulated in terms of constrained energy minimization problems. By minimizing distortions, one is able to compute, in reasonable time, simplicial maps with low energy penalty that satisfy prescribed geometric constraints. We will refer to this type of problems as *distortion minimization problems*.

Most generally the problem can be stated as follows. Given a shape S , find a transformation f of the shape S , in the family \mathcal{G} , to minimize a distortion measure E , under geometric constraints that a subset $S_0 \subset S$ of the shape maps to a given set S'_0 :

$$f^* = \operatorname{argmin}_{f \in \mathcal{G}} E(f); \tag{1}$$

$$\text{s.t. } f(S_0) = S'_0.$$

We refer to a solution f^* of problem (1) as an *optimal mapping*. The resulting problem involves a non-convex



Fig. 2 An example of surface flattening employed for texture representation of 3D faces in neural networks. From the left to right: 3D face dense reconstruction of [39] neural network, Tutte embedding into

a square, Tutte embedding into a convex shape inside a square (used in the training of [39]) and optimized Tutte embedding, obtained by minimizing isometric distortion via the projected Newton method

objective, defined over a highly non-convex domain.¹ This leads to complex non-linear optimization problems for which standard methods are not effective. Consequently, existing approaches to minimizing distortions are aimed at computing sufficiently-good minimizers of (1) with limited guaranties of obtaining a global minimum. Nevertheless, there are numerous methods for obtaining an approximate solution of the above problem, where optimization is guided by some heuristics, or by employing some indirect approaches.

Distortion minimization problem is especially prevalent in two and three dimensions, where many real world problems are encountered. This includes applications in: digital geometry processing and graphics [17,74,83,89,102,113,121] (Figs. 11, 6c), image processing [18,46] (Fig. 2), computer vision [8,51,55,75] (Figs. 2, 6d), computer-aided geometric design [32,119] (Fig. 20), physical simulations [86,116,120] (Fig. 6b) and medical imaging [12,19,43,78,91] (Figs. 1, 10 (top)).

Although there exist many approaches to the distortion minimization problem, methods used in practice lack a rigorous mathematical foundation. Many existing methods rely on heuristics built upon empirical observations. Whereas mathematical studies, often, dive deeply into an abstract theory that is far remote from practice.

In this work, we attempt to close the gap between practical applications and the underlying mathematics by exploring distortion minimization problem in depth, suitably formulating the problem with a proper balance between mathematical rigor and practical considerations. We analyze computational aspects of the existing methods and provide novel theoretical findings on how geometric distortions and methods used for minimizing these distortions behave in practice.

Our main contributions are:

- We provide a formal definition of the key concepts of the optimal mapping problem and of the related distortion measures, which so far have not been formalized in a unified manner (Sects. 2 and 3).
- We give an extensive overview of the relevant methods for distortion minimization and inversion-free mappings, and we analyze important computational aspects of these methods (Sects. 4–7). In particular, we provide a survey of recent algorithms for non-linear optimization (Sect. 4.3), we compare the continuous versus discrete problem of optimal mapping and explaining the inherent differences between discrete and continuous maps (Sects. 7).
- We identify properties of the problem that, to the best of our knowledge, have never before appeared in the literature: (i) characterization of fundamental properties of convex distortion measures and introduction of convex

geometric distortions using symmetric gauge functions (Sect. 8); (ii) the multi-resolution invariance of distortion measures (Sect. 9).

The paper is organized as follows. In Sect. 2 we focus on a continuous formulation, where we define the appropriate set of domains and transformations between them (Sect. 2.1). We then proceed to formally introduce local functionals that enable to quantify locally how much a transformation distorts the shape it operates on (Sect. 2.4). A canonical characterization of distortion measure is given in Sect. 2.5. We then address the discrete setup in Sect. 3, and formulate the distortion minimization problem for the discrete case (Sect. 4). The rest of the paper covers different optimization schemes (Sect. 4.3), analyzes convexity of the underlying minimization problem (Sect. 8) and its dependence on mesh resolution (Sect. 9). Our paper is concluded with an analysis of certain numerical aspects of the problem (Appendix A). Furthermore, we provide supplemental material with additional results, including a discussion of a variational-based formulation of the optimal mapping problem.

Since a significant part of our paper is dedicated to formalization of well-established concepts, readers who are well versed in the background and are more interested in novel results may start reading the paper from Sect. 8.

2 Continuous Problem

We first consider a general formulation of the problem for continuous, but not necessarily everywhere-differentiable, maps between Euclidean domains. We are interested in theory that includes non-differentiable maps, since later on, in discrete formulation of the problem, we will be dealing with simplicial maps which are the main object of interest in geometric processing. These maps constitute a family of piecewise linear functions that are non-differentiable on simplex faces. Formulations for non-differentiable maps, established in this section, will allow almost seamless transition between continuous and discrete scenarios. We will return to this point later in Sect. 7, after we have introduced the discrete setup. In the next section, we provide a formal definition of the relevant family of maps, \mathcal{G} , and domains over which these maps operate, then we define measure of distortion for these maps (Sect. 2.4).

2.1 Domains and Maps

We are concerned with two families of maps: continuous locally injective maps and smooth locally injective maps of Euclidean domains. The purpose of this section is to rigorously define these families of maps and outline the relevant notation that will be used throughout the rest of the paper.

¹ See Sect. 8 for the explanation of what we mean by a “highly non-convex” domain.

Given a set $S \subseteq \mathbb{R}^n$, we will denote by $\text{int}(S)$ the interior of the set. For a map of the form

$$f : S \subseteq \mathbb{R}^n \rightarrow S' \subseteq \mathbb{R}^n,$$

we will denote by $\text{Dom}(f)$ the domain of f , so $\text{Dom}(f) = S \subseteq \mathbb{R}^n$, and by $\text{Img}(f)$ we will denote the co-domain (image) of f . Throughout the paper we will call the domain of map a *source* and the co-domain a *target*.

Finally, a compact set $S \subset \mathbb{R}^n$ is called a *proper domain of \mathbb{R}^n* if it has non-empty interior, that is, $\text{int}(S) \neq \emptyset$.

Definition 2.1 (Local homeomorphism). Let S be proper domain of \mathbb{R}^n , a continuous map

$$f : S \rightarrow S' \subset \mathbb{R}^n,$$

is called a *local homeomorphism* if almost everywhere in $\text{int}(S)$ there is a neighborhood of $\mathbf{r} \in \text{int}(S)$, called a *local neighborhood of f at \mathbf{r}* , on which f is a continuous bijective map. We denote the family of such homeomorphisms by $\text{Hom}(\mathbb{R}^n)$. Wherever we will wish to restrict the family by fixing domain of the maps $\text{Dom}(f)$, or both the domain and the co-domain $\text{Img}(f)$, we will write $\text{Hom}(S, \mathbb{R}^n)$ or $\text{Hom}(S, S')$, accordingly. Thus, $f \in \text{Hom}(S, \mathbb{R}^n)$ implies that $\text{Dom}(f) = S$ and $f \in \text{Hom}(S, S')$ implies that it is also required that $\text{Img}(f) = S'$.

Definition 2.2 (Local diffeomorphism). Let S be proper domain of \mathbb{R}^n . A continuous map

$$f : S \rightarrow S' \subset \mathbb{R}^n,$$

is called a *local diffeomorphism* if almost everywhere on $\text{int}(S)$ there is a neighborhood of $\mathbf{r} \in \text{int}(S)$, called a *local neighborhood of f at \mathbf{r}* , on which f is a smooth bijective map with smooth inverse. [Note that any local diffeomorphism is a local homeomorphism but not visa versa.] We will denote the family of such diffeomorphism by $\text{Diff}(\mathbb{R}^n)$. If we want to restrict the family by fixing the proper domain S or both the proper domain and the co-domain S' we will write $\text{Diff}(S, \mathbb{R}^n)$ or $\text{Diff}(S, S')$ accordingly. So that for any $f \in \text{Diff}(S, \mathbb{R}^n)$, we have $\text{Dom}(f) = S$; and for any $f \in \text{Diff}(S, S') \subseteq \text{Diff}(S, \mathbb{R}^n)$, we have $\text{Img}(f) = S'$.

The above definitions are local in nature, and therefore we do not require that maps in $\text{Diff}(\mathbb{R}^n)$ are everywhere differentiable.

Definition 2.3 (Local first-order equivalence). Assume that $f, h \in \text{Hom}(\mathbb{R}^n)$ and $\mathbf{r}_0 \in \mathbb{R}^n$ such that $f(\mathbf{r}_0) = h(\mathbf{r}_0)$. We say that f and h are *first-order equivalent on \mathbf{r}_0* , $f \simeq h$, if there is a neighborhood $N_0 \subset \text{Dom}(f) \cap \text{Dom}(h)$ of \mathbf{r}_0 such that

$$\|f(\mathbf{r}) - h(\mathbf{r})\| = o(\|\mathbf{r} - \mathbf{r}_0\|), \forall \mathbf{r} \in N_0.$$

It is a routine procedure to verify that local first-order equivalence at \mathbf{r}_0 defines an equivalence relation on $\text{Hom}(\mathbb{R}^n)$.

As hinted by the above equivalence relation, we will be interested in only first order approximations of the maps, thereby we will be able to approximate $\text{Hom}(\mathbb{R}^n)$ by piecewise linear maps in $\text{Diff}(\mathbb{R}^n)$. The benefit of working with $\text{Diff}(\mathbb{R}^n)$, is that we can take advantage of differentiability of the latter family of maps. Singular values of Jacobians will be extremely useful when we come to analyze distortion measures, so much so that we dedicate a separate lemma to recall two crucial facts about Jacobians of local diffeomorphism and to fix the relevant notation.

Lemma 2.1 (Singular values of Jacobian). Let $f \in \text{Diff}(\mathbb{R}^n)$. We denote by $df_{\mathbf{r}}$ the Jacobian of f at $\mathbf{r} \in \text{Dom}(f)$. Then, for each $\mathbf{r} \in \text{Dom}(f)$, $df_{\mathbf{r}}$ is a full rank square matrix n -by- n and, therefore, all its singular values are positive. We will denote singular values in the descending order by $\sigma_1(df_{\mathbf{r}}), \dots, \sigma_n(df_{\mathbf{r}})$.

Proof It follows directly from the definition of $\text{Diff}(\mathbb{R}^n)$, since it implies invertibility of each $f \in \text{Diff}(\mathbb{R}^n)$. \square

2.2 Local Canonical Representation of Maps

By considering local properties of maps we provide a unified treatment of mapping,

$$f : S \subseteq \mathbb{R}^m \rightarrow S' \subseteq \mathbb{R}^d, \tag{2}$$

for any dimensions $m, d \geq 2$. In particular, as long as S is locally a manifold of dimension n , we can always transform (2) into an easier case of dimensions $n = m = d$, covered in the previous section.

Indeed, assume that any sufficiently small neighborhood $N_0 \subset S$ of \mathbf{r}_0 is a n -dimensional manifold. Then, by definition of a diffeomorphism, f embeds n -manifold $N_0 \subset \mathbb{R}^m$ onto n -manifold $f(N_0) \subset \mathbb{R}^d$. With appropriate transition maps $\phi : \mathbb{R}^m \rightarrow \mathbb{R}^n$ and $\psi : \mathbb{R}^d \rightarrow \mathbb{R}^n$, we therefore have:

$$\tilde{f}(\mathbf{r}) \triangleq [\psi \circ f \circ \phi^{-1}](\mathbf{r}), \forall \mathbf{r} \in N_0, \tag{3}$$

where \tilde{f} is of the form

$$\tilde{f} : \mathbb{R}^n \rightarrow \mathbb{R}^n.$$

We will show in Sect. 3.1 that, for triangulated domains, it is possible to pick ϕ and ψ so that their composition with f does not change the amount of distortion that f causes, and therefore we can locally substitute f by \tilde{f} without affecting the distortion values. To summarize, we can always transform f to a function \tilde{f} that locally maps \mathbb{R}^n into \mathbb{R}^n . We refer to such map \tilde{f} as a *canonical form* of a local diffeomorphism f at \mathbf{r}_0 .

2.3 Common Types of Maps

Here we give examples of different types of maps that minimize distortion measures that have intuitive interpretations. Such maps are well studied and provide a reference point against which one can bench mark and compare other maps obtained by a given distortion minimization scheme. Once we introduce these distortions, we provide the counterpart list of distortion measures that those maps minimize.

Definition 2.4 (Rigid transformation). A map $f \in \text{Diff}(\mathbb{R}^n)$ is called a *rigid transformation*, or an *isometric map* of proper domain $S = \text{Dom}(f)$, if it preserves distances.² Compositions of reflections, translations and rotations are rigid transformations.

Definition 2.5 (Harmonic maps). A map $f \in \text{Diff}(\mathbb{R}^n)$ is called *harmonic* if it is a minimizer of the Dirichlet energy functional, $\int_S \|df_r\|^2 d\mathbf{r}$. Harmonic maps are among the most studied maps in applied mathematics and functional analysis.

Definition 2.6 (Conformal maps). A map $f \in \text{Diff}(\mathbb{R}^n)$ is called *conformal* map of a proper domain $S = \text{Dom}(f)$, if for each point $\mathbf{r} \in \text{int}(S)$ it scales the space uniformly in every direction. This can be stated formally as

$$\|df_r \cdot \hat{u}_1\|_2 = \|df_r \cdot \hat{u}_2\|_2, \tag{4}$$

where $\hat{u}_1, \hat{u}_2 \in \mathbb{R}^n$ denote two unit vectors. According to the above notation, a conformal map $f(\mathbf{r})$ is isometric in S if

$$\forall \mathbf{r} \in \text{int}(S) : |\det df_r| = 1. \tag{5}$$

Intuitively, conformal maps are angle-preserving maps.

Definition 2.7 (Equi-volume maps). Property (5) in Definition 2.6 by itself defines a class of *equi-volume transformations*. As the terminology suggests, these maps preserve volume.

Definition 2.8 (Quasi-conformal maps). A map $f \in \text{Diff}(\mathbb{R}^n)$ is called *quasi-conformal* if there exist $K \in [1, \infty)$ such that for any unit vectors \hat{u}_1, \hat{u}_2 and $\mathbf{r} \in \text{Dom}(f)$,


$$\frac{1}{K} < \frac{\|df_r \cdot \hat{u}_1\|_2}{\|df_r \cdot \hat{u}_2\|_2} < K. \tag{6}$$

Definition 2.9 (Quasi-isometry maps). A map $f \in \text{Hom}(\mathbb{R}^n)$ is called *quasi-isometric* if there exist a number $C \in [1, \infty)$ such that for each $\mathbf{r}_1, \mathbf{r}_2 \in \text{Dom}(f)$,

$$\frac{1}{C} \|\mathbf{r}_1 - \mathbf{r}_2\| \leq \|f(\mathbf{r}_1) - f(\mathbf{r}_2)\| \leq C \|\mathbf{r}_1 - \mathbf{r}_2\|. \tag{7}$$

² We always assume Euclidean metric.

These classes of maps are rich and well studied objects, whose theoretical understanding is based on various areas of mathematics and many deep mathematical insights, that combine topology, algebra and more. Studying these maps in detail is, of course, beyond the scope of our work. Nevertheless, we will often invoke these definitions, whenever it will be important to make a distinction that a given optimization method might converge to one type of functions and not to the another. For example, while, according to the Riemann mapping theorem, there exist an abundance of continuous conformal maps in two dimensions, higher dimensional domains can be mapped only *quasi-conformally* [78]. Mutual relations between these classes of maps are summarized in the diagram below:



Having provided the set of maps that we are interested in, and providing examples of the most important types of maps, we move on to introduce distortions that measure how a map distorts locally its domain.

2.4 Distortions

Similarly to local definitions of $\text{Hom}(\mathbb{R}^n)$ and $\text{Diff}(\mathbb{R}^n)$ we are interested in local definition of a distortions—the resulting framework will be applicable to maps that have irregular points, such as non-differentiable points, and it will therefore apply for simplicial maps as well.

Definition 2.10 (Local functional). Adopting the notation of Sect. 2.1, we define a *local functional* as a map

$$\mathcal{D} : \{(f, \mathbf{r}) \mid f \in \text{Hom}(\mathbb{R}^n), \mathbf{r} \in \text{Dom}(f)\} \rightarrow \mathbb{R}. \tag{8}$$

That is, \mathcal{D} maps a pair (f, \mathbf{r}) ($f \in \text{Hom}(\mathbb{R}^n)$, $\mathbf{r} \in \text{Dom}(f)$), to a real number. For fixed f , \mathcal{D} is a map from $\text{Dom}(f)$ to \mathbb{R} , and for fixed $\mathbf{r} \in \mathbb{R}^n$, \mathcal{D} is a functional on $\{f \in \text{Hom}(\mathbb{R}^n) \mid \mathbf{r} \in \text{Dom}(f)\}$.

A local functional provide us with the basis for the definition of a distortion:

Definition 2.11 (Distortion). *Distortion* is a local functional that satisfies the following properties:

1. **Coordinate frame invariance.** Distortion measures, used in geometric processing, are motivated by some physical quantities and therefore, are expected to be independent of a specific orthogonal coordinates selected to represent the source and target domains. Consequently,

distortion measures need to be invariant to composition of f with rigid transformations. In other words, if $f \in \text{Hom}(S, S')$, and R_1 is a rigid transformation of S , while R_2 is rigid a transformation of S' (see Definition 2.4), then for each $y \in R_1(S)$,

$$\mathcal{D}(R_2 \circ f \circ R_1^{-1}, y) = \mathcal{D}(f, R_1^{-1}(y)). \tag{9}$$

2. **First-order precision.** Assume that $f, h \in \text{Hom}(\mathbb{R}^n)$ are first-order equivalent on $\mathbf{r}_0 \in \text{Dom}(f) \cap \text{Dom}(h)$ in the sense of Definition 2.3, then

$$\mathcal{D}(f, \mathbf{r}_0) = \mathcal{D}(h, \mathbf{r}_0).$$

The above definitions are based on minimal requirement that a local functional has to satisfy to define a distortion. However, it is often desirable to impose additional regularity conditions on a distortion. Distortion that satisfy all these additional requirements are called *regular*.

Definition 2.12 (Regular distortion). A *regular distortion* \mathcal{D} is a distortion that in addition satisfies:

1. **Normalization.** Denote by I_S an identity map of a set S . Distortion $\mathcal{D}(f, \mathbf{r})$ is called *normalized* if: **(i)** $\mathcal{D}(f, \mathbf{r}) \in [\zeta, \infty)$ for $\zeta \geq 0$ and any $f \in \text{Diff}(\mathbb{R}^n), \mathbf{r} \in \text{Dom}(f)$; **(ii)** the following conditions are met:

$$\mathcal{D}(I_{\text{Dom}(f)}, \mathbf{r}) = \zeta, \quad \forall \mathbf{r} \in \text{Dom}(f). \tag{10}$$

In particular, for all rigid transformations R on $\text{Dom}(f)$,

$$\mathcal{D}(R, \mathbf{r}) = \zeta, \quad \forall \mathbf{r} \in \text{Dom}(f). \tag{11}$$

Usually, the value of ζ is set to $\zeta = 0$ or $\zeta = 1$.

2. **Symmetry under inversions.** We denote by $f_{\mathbf{r}_0}^\dagger \in \text{Hom}(\mathbb{R}^n)$ a *local inversion* of map f at $\mathbf{r}_0 \in \text{Dom}(f)$, if there exist a neighborhood N of $\text{Dom}(f)$ at \mathbf{r}_0 , such that:

$$\mathbf{r} = [f_{\mathbf{r}_0}^\dagger \circ f](\mathbf{r}), \quad \forall \mathbf{r} \in N.$$

We call \mathcal{D} *symmetric*, if for any $f \in \text{Diff}(\mathbb{R}^n), \mathbf{r} \in \text{Dom}(f)$, and any local inversion $f_{\mathbf{r}}^\dagger$ of f at \mathbf{r} ,

$$\mathcal{D}(f, \mathbf{r}) = \mathcal{D}(f_{\mathbf{r}}^\dagger, f(\mathbf{r})). \tag{12}$$

In other words, symmetric distortions do not distinguish between switching the role of the source and the target, i.e., they assign the same distortion when deforming the source into the target or when doing the inverse (i.e., deforming the target back into the source). For this reason, symmetric distortions have been used extensively in many computer graphics applications [49,89,92,95].

3. **Bottom barrier property.** A sequence of maps $f^{(j)} \in \text{Diff}(\mathbb{R}^n), j = 1, 2, \dots$, is called a *bottom barrier sequence* at \mathbf{r} , if there exist vectors $u_1, u_2 \in \mathbb{R}^n$ and a number $Q < \infty$ so that, for each $j, f^{(j)}$ is differentiable at \mathbf{r} and

$$\|df_{\mathbf{r}}^{(j)} u_1^\top\|_2 < Q, \quad \lim_{j \rightarrow \infty} \|df_{\mathbf{r}}^{(j)} u_2^\top\|_2 = \infty, \tag{13}$$

where $df_{\mathbf{r}}^{(j)}$ denotes the Jacobian of $df^{(j)}$ at \mathbf{r} . We say that \mathcal{D} has the *bottom barrier property* if

$$\lim_{j \rightarrow \infty} \mathcal{D}(f^{(j)}, \mathbf{r}) = \infty,$$

for any bottom barrier sequence $\{f^{(j)}\}_{j=1}^\infty$ at \mathbf{r} .

4. **Top barrier property.** A sequence $f^{(j)} \in \text{Diff}(\mathbb{R}^n), j = 1, 2, \dots$ is called a *top barrier sequence* at \mathbf{r} if there exist vectors $u_1, u_2 \in \mathbb{R}^n$ and a number $\varepsilon > 0$ such that, for each $j, f^{(j)}$ is differentiable at \mathbf{r} and

$$\|df_{\mathbf{r}}^{(j)} u_1^\top\|_2 > \varepsilon, \quad \lim_{j \rightarrow \infty} \|df_{\mathbf{r}}^{(j)} u_2^\top\|_2 = 0. \tag{14}$$

We say that \mathcal{D} has the *top barrier property* if

$$\lim_{j \rightarrow \infty} \mathcal{D}(f^{(j)}, \mathbf{r}) = \infty$$

for any top barrier sequence $\{f^{(j)}\}_{j=1}^\infty$ at \mathbf{r} .

5. **Smoothness almost everywhere.** Intuitively, we expect a small discrepancy in $\mathcal{D}(f, \mathbf{r})$ when either f or \mathbf{r} are slightly changed. To formalize this intuitive property we use the *weak derivative* of \mathcal{D} at \mathbf{r}

$$\frac{\partial \mathcal{D}(f, \cdot)}{\partial \Delta} \Big|_{\mathbf{r}} \triangleq \limsup_{\varepsilon \rightarrow 0} \frac{\mathcal{D}(f, \mathbf{r} + \varepsilon \Delta) - \mathcal{D}(f, \mathbf{r})}{\varepsilon \|\Delta\|}, \tag{15}$$

and define special derivative of distortion with respect to deformation f

$$\frac{\partial \mathcal{D}(\cdot, \mathbf{r})}{\partial A} \Big|_f \triangleq \limsup_{\varepsilon \rightarrow 0} \frac{\mathcal{D}(f + \varepsilon A, \mathbf{r}) - \mathcal{D}(f, \mathbf{r})}{\varepsilon \|A\|}, \tag{16}$$

where A is a linear map in $\text{Diff}(\mathbb{R}^n)$. Distortion that satisfies these two properties is referred to as *smooth almost everywhere (a.e.)*, or to as *weakly differentiable*.

2.5 Canonical Representation of Distortions

With Definition 2.11 in place, we are now in the position to introduce canonical representations for distortions. Such representations constitute a crucial step in analyzing distortions, as it provides a very convenient way to characterize different distortions. Moreover, when not written in their canonical

form, the arguments of distortions are prone to contain extra degrees of freedom. This might lead to situations where the same distortion can be represented in multiple inherently different ways, leading to certain difficulties in their processing. Canonical representation avoids the unnecessary ambiguity in distortion representations and allows to treat all distortions in a unified way. To obtain the canonical representations we rely on basic linear algebra properties of Jacobian df_r of $f \in \text{Diff}(\mathbb{R}^n)$ at $r \in \text{Dom}(f)$, stated in Lemma 2.1. We characterize the distortion by means of the following fundamental theorem [81,89]:

Theorem 2.1 (Canonical representation of distortions). *Let $f \in \text{Diff}(\mathbb{R}^n)$ and $r \in \text{Dom}(f)$, then local functional \mathcal{D} of the form (8), is a distortion (according to Definition 2.11) iff it can be expressed as a function of singular values of the Jacobian df_r . That is,*

$$\mathcal{D}(f, r) = \tilde{\mathcal{D}}(\sigma_1(df_r), \dots, \sigma_n(df_r)) \tag{17}$$

for a map $\tilde{\mathcal{D}} : \mathbb{L}^n \rightarrow \mathbb{R}$, where \mathbb{L}^n is the lower half-space of \mathbb{R}^n located below the main diagonal, i.e.,

$$\mathbb{L}^n \triangleq \{r \in \mathbb{R}^n \mid r_1 \geq r_2 \geq \dots \geq r_n > 0\}. \tag{18}$$

We call such a representation the canonical representation.

Proof Let \mathcal{D} be a distortion according to Definition 2.11. We first show that $\mathcal{D}(f, r_0)$ is necessarily a function of the entries of the Jacobian matrix of f at r_0 . By first property (9) in Definition 2.11, we can always pick appropriate rigid transformations R_1 and R_2 that rotate and shift $\text{Dom}(f)$, $\text{Img}(f) \subset \mathbb{R}^n$ so that both $r_0 \in \text{Dom}(f)$ and $f(r_0) \in \text{Img}(f)$ are moved to the origin. Therefore, without loss of generality we assume that

$$r_0 = f(r_0) = \underbrace{(0, 0, \dots, 0)}_n.$$

Second, let N_0 be a sufficiently small local neighborhood of \mathbb{R}^n at r_0 . Let df_0 be the Jacobian of f at 0 and denote by $df_0(r)$ the linear map

$$df_0(r) : r \mapsto (df_0)r, \quad r \in N_0.$$

Since f is by assumption a diffeomorphism on N_0 , it can be linearly approximated by the first term in its Taylor series expansion. Hence, $df_0(r)$ and f are first-order equivalent on r_0 , (see Definition 2.3). Therefore, it follows from Definition 2.11 that

$$\mathcal{D}(f, 0) = \mathcal{D}(df_0, 0).$$

Consequently, $\mathcal{D}(f, 0)$ is a function of the entries of df_0 , the Jacobian matrix of f at r_0 . Finally, let $df_0 = U\Sigma V^T$ be the SVD of the Jacobian, so that

$$\Sigma = \text{diag}(\sigma_1(df_0), \dots, \sigma_n(df_0)),$$

where the singular values $\sigma_1(df_0), \dots, \sigma_n(df_0)$ are in the descending order. Then, applying property (9) of Definition 2.11 with $R_1 = V$, $R_2 = U^T$ yields

$$\mathcal{D}(df_0(r), 0) = \tilde{\mathcal{D}}(\sigma_1(df_0), \dots, \sigma_n(df_0)),$$

for the corresponding function $\tilde{\mathcal{D}} : \mathbb{L}^n \rightarrow \mathbb{R}$.

To prove the other direction, note that by definition the local functional $\tilde{\mathcal{D}} : \mathbb{L}^n \rightarrow \mathbb{R}$ is operating on $\sigma_1(df_r), \dots, \sigma_n(df_r)$, thus $\tilde{\mathcal{D}}$ satisfies the first-order precision property. Indeed, if two maps are first-order equivalent on r this implies that they have the same Jacobian on r . Thus, it remains to show that $\tilde{\mathcal{D}}$ satisfies the coordinate frame invariance. The latter follows from the fact that, for rigid transformations R_1 and R_2 , we have:

$$\sigma_i(d(R_2 \circ f \circ R_1^{-1})_r) = \sigma_i(df_{R_1^{-1}(r)}).$$

The above equality completes the proof. □

Hereinafter we will drop the distinction between $\tilde{\mathcal{D}}$ and \mathcal{D} , and will occasionally write $\mathcal{D}(\sigma_1(df), \dots, \sigma_n(df))$, $\mathcal{D}(\sigma_1, \dots, \sigma_n)$, or $\mathcal{D}(\Sigma)$ — they all represent the same function, defined in (17). Expressing distortion in terms of singular values of Jacobian establishes a differential definition of distortions. The advantage of a differential definition is that it explicitly factors in the first-order precision and coordinate invariance, leaving out nuisance parameters and retaining only the essential n degrees of freedom.

Remark 2.1 It is common in geometry processing to represent distortions by singular values [49]. Theorem 2.1 is similar to propositions presented in [81,89]. In particular, Rabinovich et al. [89] have proven a variation of Theorem 2.1 which shows that rotation-invariant geometric measures can be represented by the signed SVD, namely by using the decomposition $df_r = U\tilde{\Sigma}V^T$, where U and V are positive orthonormal matrices, and $\tilde{\Sigma}$ is an arbitrary diagonal matrix. Unlike this work, our paper uses the unsigned SVD, leading to a slightly different formulation in which the signs of $\det df_r$ are prescribed by a set of separate orientation constraints (for details, see Sect. 5).

On the one hand, due to first order precision only the Jacobian of a map matters. On the other hand, the distortion is invariant to rigid transformations, so for fixed $r \in \mathbb{R}^n$ distortions are, in fact, functionals defined over a quotient space of

linear maps in \mathbb{R}^n , with equivalence given by unitary transformations, i.e., two maps

$A(\mathbf{r}) : \mathbf{r} \mapsto A\mathbf{r}$ and $B(\mathbf{r}) : \mathbf{r} \mapsto B\mathbf{r}$, $A, B \in \mathbb{R}^{n \times n}$, $\mathbf{r} \in \mathbb{R}^n$, are equivalent if there are unitary matrices $R_1, R_2 \in \mathbb{R}^{n \times n}$ such that $A = R_2 B R_1^T$. We denote the above equivalence by

$$A(\mathbf{r}) \sim B(\mathbf{r}), \tag{19}$$

or by $A \sim B$, for short. Clearly, ‘ \sim ’ is an equivalence relation and if $[df_{\mathbf{r}}]_{\sim}$ is an equivalence class of $df_{\mathbf{r}}$ with respect to ‘ \sim ’, then, according to (9),

$$\mathcal{D}(f, \mathbf{r}) = \mathcal{D}(h), \quad \forall h \in [df_{\mathbf{r}}]_{\sim}.$$

Furthermore, SVD of $df_{\mathbf{r}}$,

$$df_{\mathbf{r}} = U \text{diag}(\sigma_1(df_{\mathbf{r}}), \dots, \sigma_n(df_{\mathbf{r}})) V^T, \tag{20}$$

suggests a convenient choice of local coordinate bases

$$U = [u_1, \dots, u_n] \text{ and } V = [v_1, \dots, v_n],$$

so that right-singular vectors (V) of $df_{\mathbf{r}}$ are used as basis for the neighborhood of $\mathbf{r} \in \text{Dom}(f) \subset \mathbb{R}^n$; and left-singular vectors (U) are used as a basis of a neighborhood of $f(\mathbf{r}) \in \text{Img}(f) \subset \mathbb{R}^n$ in the co-domain.

Before concluding this section, we formulate the following corollary of Theorem 2.1 that vastly simplifies mathematical characterization of regular distortion:

Corollary 2.1 *Let \mathcal{D} be a distortion measure such that $\mathcal{D}(f, \mathbf{r}) \in [\zeta, \infty)$ for $\zeta \geq 0$ and any $f \in \text{Diff}(\mathbb{R}^n)$, $\mathbf{r} \in \text{Dom}(f)$. Then, \mathcal{D} is a regular distortion in the sense of Definition 2.12 iff the canonical representation of \mathcal{D} , as a function of $(\sigma_1, \dots, \sigma_n) \in \mathbb{L}^n$, satisfies:*

1. **Normalization property:**

$$\mathcal{D}(1, 1, \dots, 1) = \zeta. \tag{21}$$

2. **Symmetry property:**

$$\mathcal{D}(\sigma_1, \dots, \sigma_n) = \mathcal{D}\left(\frac{1}{\sigma_n}, \dots, \frac{1}{\sigma_1}\right). \tag{22}$$

3. **Bottom barrier property:** when $\sigma_1 > \varepsilon > 0$,

$$\mathcal{D}(\sigma_1, \dots, \sigma_n) \rightarrow \infty \text{ as } \sigma_n \rightarrow 0. \tag{23}$$

4. **Top barrier property:** when $\sigma_n < Q < \infty$,

$$\mathcal{D}(\sigma_1, \dots, \sigma_n) \rightarrow \infty \text{ as } \sigma_1 \rightarrow \infty. \tag{24}$$

5. **Smoothness almost everywhere:** $\mathcal{D}(\sigma_1, \dots, \sigma_n)$ is a weakly differentiable map (see (15) and (16)).

Proof If R is a rigid transformation of a proper domain S , then by Definition 2.11 and Theorem 2.1,

$$\mathcal{D}(R, \mathbf{r}) = \mathcal{D}(R^{-1} \circ R, \mathbf{r}) = \mathcal{D}(I_S, \mathbf{r}) = \mathcal{D}(1, \dots, 1).$$

Therefore, conditions (21), (11) and (10) are equivalent normalization properties. The equivalence of properties (22) and (12) follows from the fact that if $\sigma_1, \dots, \sigma_n$ are singular values of a full rank $df_{\mathbf{r}}$, then $\sigma_n^{-1}, \dots, \sigma_1^{-1}$ are singular values of $df_{\mathbf{r}}^\dagger$ (and visa versa). The bottom and top barrier properties of Definition 2.12 are equivalent to (23) and (24), respectively, since multiplying vector by unitary matrix does not change its 2-norm. \square

Clearly, it is much easier to define regular distortion by means of canonical representation. In the next section, we will utilize this representation to introduce a few examples of distortions that are used in practice.

2.6 Common Distortion Measures

We proceed by providing examples of different types of distortions that are considered to be useful in practice. It should be clear from our previous discussions (see proof of Theorem 2.1), that properly normalized distortions $\mathcal{D}(f, \mathbf{r})$ are, in fact, estimates of f ’s rigidity at \mathbf{r} ; meaning they are a measure of how ”close” f is to a rigid transformation on $\text{Dom}(f)$.

As we have seen in Sect. 2.3, there are four major classes of ”nice” maps: length-preserving, harmonic, angle-preserving and volume-preserving. Such maps can be related to distortion measures which, in a sense, measure to what extend a map differs from each of the above classes.

Definition 2.13 (Harmonic distortion). The following measure, called *Harmonic distortion* or *Dirichlet energy* [45], is closely related to harmonic maps; it measures by how much a given map $f \in \text{Diff}(\mathbb{R}^n)$ stretches a small neighborhood of $\mathbf{r} \in \text{Dom}(f)$ and is defined by

$$\mathcal{D}_{\text{Dirichlet}}(f, \mathbf{r}) \triangleq \|df_{\mathbf{r}}\|_{\text{Fro}}^2 = \sum_{i=1}^n \sigma_i^2, \tag{25}$$

where $\|\cdot\|_{\text{Fro}}$ is the Frobenius norm. This distortion, is also sometimes referred to as *smoothness energy*. It is widely employed in construction of harmonic surface parameterizations [49]. Although there exist a few approaches that employ harmonic distortions for computing volumetric mappings [65,114], most of the methods based on Dirichlet energy are focused on planes and two dimensional surfaces embedded in \mathbb{R}^3 .

Definition 2.14 (Conformal distortions). A $f \in \text{Diff}(\mathbb{R}^n)$ is conformal at $\mathbf{r} \in \text{Dom}(f)$ iff

$$\sigma_i(df_{\mathbf{r}}) = \sigma_j(df_{\mathbf{r}}), \text{ for } 1 \leq i, j \leq n.$$

Therefore, *conformal distortions* quantify how much singular values deviate one from the other. Most commonly, this deviation is measured by:

- The *MIPS_{2D} distortion* [45] defined on $\text{Diff}(\mathbb{R}^2)$ as

$$\text{MIPS}_{2D}(f, \mathbf{r}) \triangleq \frac{\sigma_1}{\sigma_2} + \frac{\sigma_2}{\sigma_1} = \frac{\sigma_1^2 + \sigma_2^2}{\sigma_1\sigma_2}. \tag{26}$$

MIPS_{2D} distortion is referred to as “most isometric parameterizations”. Despite the terminology, this distortion only estimates the deviation of σ_1 from σ_2 , and thus is a metric of conformal distortion (see the related discussion in Sect. 7).

The distortion was extensively employed in early geometric processing applications, since for the discrete setup it yields convex optimization in a single vertex of a simplicial complex if all other vertices in complex are kept fixed. (We will return to this point in Sect. 8 after we introduce discrete setup.)

- MIPS_{2D} can be extended to n -dimensions as n -times the ratio between arithmetic and geometric means of squared singular values, yielding the *MIPS_{nD} distortion* [37, 81]:

$$\text{MIPS}_{nD} \triangleq \frac{\sigma_1^2 + \dots + \sigma_n^2}{(\sigma_1 \dots \sigma_n)^{2/n}} = \frac{\text{trace}((df_{\mathbf{r}})^T df_{\mathbf{r}})}{|\det(df_{\mathbf{r}})|^{2/n}}, \tag{27}$$

which, according to (25), has the following relation to the Dirichlet energy

$$\text{MIPS}_{nD}(f, \mathbf{r}) = \frac{\mathcal{D}_{\text{Dirichlet}}(f, \mathbf{r})}{|\det(df_{\mathbf{r}})|^{2/n}}. \tag{28}$$

- *Condition number distortion*, also called *linear dilatation*, is a simple and natural measure for assessing the conformality of a linear function as a ratio of its maximal and minimal singular values. Thus, it induces the following distortion of a smooth deformation

$$\mathcal{D}_{\text{conf}}(f, \mathbf{r}) \triangleq \frac{\sigma_1}{\sigma_n}. \tag{29}$$

Condition number (29) is extensively employed in various studies on geometric optimization [3,35,55,66,89, 102]. In particular, convexification algorithms (e.g., the algorithm of [55], listed in Sect. 4.3) use the condition number for estimating conformal distortion, since $\mathcal{D}_{\text{conf}}(\sigma_1, \dots, \sigma_n)$ is a quasi-convex function of singular values.

- *Quasi-conformal (qc) dilatation* is another geometric measure, employed in the classical theory of quasi-conformal maps for estimating the maximal conformal distortion induced by homeomorphic deformations. Originally, this quantity was defined by means of an abstract measure over curve families, called *modulus* [110]. The density of the qc-dilatation can be expressed by

$$\mathcal{D}_K(f, \mathbf{r}) \triangleq \max \left\{ \underbrace{\frac{\sigma_1 \dots \sigma_{n-1}}{\sigma_n^{n-1}}}_{K_I}, \underbrace{\frac{\sigma_1^{n-1}}{\sigma_2 \dots \sigma_n}}_{K_O} \right\}, \tag{30}$$

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

Quasi-conformal dilatations are often employed in mathematical analysis of quasi-conformal maps; notably they are useful for estimation of geometry-dependent bounds of conformal distortions [15,79,110].

Definition 2.15 (Volume distortions). If $f \in \text{Hom}(\mathbb{R}^n)$ is equi-volume map in the vicinity of \mathbf{r} , then

$$|\det(df_{\mathbf{r}})| = \sigma_1\sigma_2 \dots \sigma_n = 1.$$

Hence, for estimation of the volume distortion (or the area distortion in two dimensions) we employ the following measure [29,81]:

$$\mathcal{D}_{\text{vol}}(f, \mathbf{r}) \triangleq \max \left\{ |\det(df_{\mathbf{r}})|, |\det(df_{\mathbf{r}})|^{-1} \right\}, \tag{31}$$

where $|\det(df_{\mathbf{r}})|$ and $|\det(df_{\mathbf{r}})|^{-1}$ can be interpreted as assessments of the dilatation and of the compression of a local volume, accordingly. However, (31) and other volume-based measures do not satisfy the barrier properties (23) and (24) and therefore, can lead to non-desirable results during optimizations. For instance, the value of $\mathcal{D}_{\text{vol}}(f, \mathbf{r})$ can be the same for regular and nearly collapsed simplices. As a result, iterative descent algorithms for minimizing $\mathcal{D}_{\text{vol}}(f, \mathbf{r})$ can produce degenerate³ simplices, leading to adverse numerical issues. Thus, measures of volume are most often used in a combination with other distortions for improving numerical stability of optimization process (for an example of such measure see (36)).

Not surprising, the minimization of \mathcal{D}_{vol} is linked to the problems of finding volume-preserving mapping and to the closely related problem of the optimal mass transportation

³ We explain the notion of degenerate and inverted simplices in Sect. 3.

[71,96,124]. We will return to these problems at the end of Sect. 4.3.

Definition 2.16 (Isometric distortions). Isometric distortions are direct measures of the rigidity. Since singular values of an isometry all equal 1, these distortions assess the deviation of $(\sigma_1, \dots, \sigma_n)$ from the vector $(1, \dots, 1)$.

- Arguably, the most popular measure of isometric distortion, employed in geometry processing application, is *symmetric Dirichlet energy* [17,74,89,102,103,121]

$$\begin{aligned} \mathcal{D}_{SD}(f, \mathbf{r}) &\triangleq \sum_{i=1}^n (\sigma_i^2 + \sigma_i^{-2}) \\ &= \|df_{\mathbf{r}}\|_{\text{Fro}}^2 + \|df_{\mathbf{r}}^{-1}\|_{\text{Fro}}^2 \\ &= \mathcal{D}_{\text{Dirichlet}}(f, \mathbf{r}) + \mathcal{D}_{\text{Dirichlet}}(f_{\mathbf{r}}^{\dagger}, \mathbf{r}). \end{aligned} \tag{32}$$

Symmetric Dirichlet energy is a regular distortion (Definition 2.12, Corollary 2.1), and thus it contains the barrier term that prevents simplex inversions in iterative optimization algorithms (we will discuss this property in Sect. 5).

- *As-rigid-as-possible-distortion* (ARAP) is another popular isometric distortion, employed in computer graphics for surface parameterization and shape deformation [90]

$$\mathcal{D}_{\text{ARAP}}(f, \mathbf{r}) \triangleq \sum_{i=1}^n (\sigma_i - 1)^2. \tag{33}$$

Unlike symmetric Dirichlet energy, ARAP is a non-symmetric and non-barrier distortion, and thus cannot guarantee inversion-free mapping for standard algorithms in distortion minimization. For this reason, ARAP energy is often modified, by adding an inversion barrier term⁴ that prevents simplex inversions

$$\mathcal{B}(f, \mathbf{r}) \triangleq \begin{cases} \infty & \det(df_{\mathbf{r}}) \leq 0, \\ 0 & \text{else,} \end{cases} \tag{34}$$

or ARAP energy is substituted by its symmetric variant [102]

$$\mathcal{D}_{\text{SARAP}}(f, \mathbf{r}) \triangleq (\sigma_1 - 1)^2 + (\sigma_n^{-1} - 1)^2. \tag{35}$$

- Fu et al. [37] have introduced the family of *advanced MIPS* distortions (APIMS) as variants of MIPS distortions (26) and (27), modified for assessing isometric

⁴ By distortions with barrier terms we refer to measures $\mathcal{D}(f, \mathbf{r}) + \mathcal{B}(f, \mathbf{r})$, where \mathcal{D} is a first-order distortion and \mathcal{B} is defined by (34). As explained in Sect. 8, measures $\mathcal{D} + \mathcal{B}$ extend the essential properties of Definition 2.11 to the domain $\mathbb{R} \cup \{\infty\}$.




Fig. 3 Distribution of distortion obtained in the flattenings of the face model from Fig. 2. Mesh colors encode the following distortions (from the top to the bottom): (32), (29) and (25)

distortions. AMIPS includes, among others, the following type of distortions:

$$\mathcal{D}_{\text{AMIPS}}^{nD}(f, \mathbf{r}) \triangleq \exp(\text{MIPS}_{nD}(f, \mathbf{r}) + 0.5 \cdot \mathcal{D}_{\text{vol}}(f, \mathbf{r})). \tag{36}$$

Because of the exponent, the distortion $\mathcal{D}_{\text{AMIPS}}$ grows faster on barrier sequences than other isometric distortions. This makes (36) an attractive measure for assessing geometric distortions in applications that are particularly sensitive to deformations with ill-conditioned Jacobians. For instance, $\mathcal{D}_{\text{AMIPS}}$ is employed in recent methods for generating tetrahedral meshes to avoid poorly-shaped simplices [50,52].

- In the classical geometric analysis, the rigidity is often measured by the, so-called, *Quasi-isometric (qi) dilatation*, which gives rise to a symmetric distortion that is closely related to the notation of *qi-mappings*, introduced in Definition 2.9,⁵

$$\mathcal{D}_{\text{iso}}(f, \mathbf{r}) \triangleq \max\{\sigma_1, \sigma_n^{-1}\}. \tag{37}$$

This measure is used in both theoretical studies [15] and in practice [56,80,95], where it is often normalized according to the relative sizes of the source and target domains.

Fig. 3 illustrates a few distortion measures obtained by mapping a triangulated surface onto different planar domains.

⁵ If f is differentiable *qi*-mapping in neighborhood N of \mathbf{r} , then $\mathcal{D}_{\text{iso}}(f, \mathbf{r})$ is the infimum over numbers C that satisfy (7) in N .




Fig. 4 We compute surface parameterization by minimizing the following measures (from the top to the bottom): isometric distortion \mathcal{D}_{SD} , (32); conformal distortion \mathcal{D}_{conf} , (29); the linear combination of the two distortions $\mathcal{D}_\lambda = (1 - \lambda) \mathcal{D}_{conf} + \lambda \mathcal{D}_{SD}$, for $\lambda = 10^{-3}$. We cut the brain hemisphere surface along the selected edges (highlighted in green) into meshes with disc-like topology, and then initialize the problem with Tutte embedding. Fixed boundary parameterizations are visualized from the left to the right, as follows: we show the ‘outer’ and the ‘inner’ sides of textured surfaces, and target meshes, obtained after 30 iterations of BCQN solver. Minimizing \mathcal{D}_{SD} yields high angle distortions around the cut, whereas minimizing \mathcal{D}_{conf} causes very strong shrinking of edges on the outer side of the hemisphere. At the same time, constrained parameterization with \mathcal{D}_λ attains both low conformal and low isometric distortions. Thus, compared with other two measures, minimizing \mathcal{D}_λ causes less visual artifacts in texture mapping

Before we complete this section, we would like to note that a new regular distortion can be derived by applying simple arithmetic operations on existing regular distortions. In particular, it is easy to show that if \mathcal{D}_1 and \mathcal{D}_2 are regular, then

$$\mathcal{D}_\lambda = \lambda \mathcal{D}_1 + (1 - \lambda) \mathcal{D}_2, \quad \lambda \in [0, 1], \tag{38}$$

and $\exp(\mathcal{D}_\lambda)$ are also regular. Distortions \mathcal{D}_λ and $\exp(\mathcal{D}_\lambda)$ generated this way can be quite useful: they can be used to mitigate drawbacks of the original distortions and to craft a new distortion for application-specific tasks. For example, it could be too restrictive to compute “as-length-preserving-as-possible” maps in cases that involve deformations of shapes with complex geometry and under hard positional constraints. As illustrated by Fig. 4, instead of minimizing the standard isometric distortion, it might be better to minimize a more flexible combination of regular distortions, defined according to (38). Likewise, if \mathcal{D} is regular, then $\exp(\mathcal{D})$ is also regular, but has a ‘stronger’ barrier term than \mathcal{D} . Thus, minimizers of $\exp(\mathcal{D})$ are less likely to yield simplicial maps

with ill-conditioned Jacobians, so that using $\exp(\mathcal{D})$ instead of \mathcal{D} can lead to more stable numerical computations [37].

This concludes our exposition of geometric distortions. In Sect. 8 we will return to distortion measures and discuss methods for designing convex distortions.

3 Discrete Problem

We now use concepts introduced in Sect. 2 to reformulate the distortion minimization problem in a more practical form, that will serve us in the rest of the paper.

We assume that triangulated domains, considered in our paper, are manifold meshes, represented by simplicial complexes, and that mappings between these domains are piecewise affine functions, represented by simplicial maps. That is, in 2D and 3D, by “simplicial complexes” we refer to triangular and tetrahedral meshes, and by “simplicial maps” we refer to piecewise affine transformations of meshes.

Although our assumption implies that simplicial maps are continuous functions, we refer to the simplicial mapping problem as to a “discrete problem” because a simplicial complex can be represented by a finite number of entries that encode vertex positions and describe how vertices are connected to form the complex simplices.

Let \mathcal{V} be a *vertex* set and let \mathcal{S} be a consistently oriented *simplex* set of these vertices. Denote by $\mathbf{M} = (\mathcal{S}, \mathcal{V}, \mathbf{y})$ a simplicial complex of \mathcal{V} and \mathcal{S} embedded in \mathbb{R}^m in such a way that $\mathbf{y} \in \mathbb{R}^{m|\mathcal{V}|\times 1}$ is the column stack of vertex coordinates in \mathbb{R}^m . Further, denote by $\dim(\mathbf{M}) = n$ that all simplices in \mathcal{S} are n -dimensional.⁶ Denote by $\text{conv}(s)$ the closed convex hull of simplex $s \in \mathcal{S}$, embedded in \mathbb{R}^m according to the coordinates specified in \mathbf{y} (*source coordinates*).

We assume that interiors of $\text{conv}(s)$ are disjoint for different simplices and we define

$$\text{conv}(\mathbf{M}) \triangleq \bigcup_{s \in \mathcal{S}} \text{conv}(s).$$

A *simplicial map* f of \mathbf{M} is then a piecewise affine function

$$f : \text{conv}(\mathbf{M}) \mapsto \mathbb{R}^d, \tag{39}$$

where $n \leq d$ and by the “piecewise affine” we means that the restriction $f_s \triangleq f|_{\text{conv}(s)}$ is an affine map for each simplex s .

We assume that any simplex $s \in \mathcal{S}$ can be represented as a $(n + 1)$ -tuple $(v_1 \dots, v_{n+1})$ of vertices that constitute s , where the order of vertices reflects the simplex orientation; we write $v_j \in s$ to denote that simplex s is built on these

⁶ For example, $\dim(\mathbf{M}) = 2$ if \mathcal{S} are triangles and $\dim(\mathbf{M}) = 3$ if \mathcal{S} are tetrahedra.




Fig. 5 Illustrating equivalent representations of simplicial maps. We choose a map f of the triangle mesh M and show (from left to right): piecewise affine representation $f = f[\mathbf{x}]$, defined according to (40); a discrete representation of f as a function $f_{\mathcal{V}} = f|_{\mathcal{V}}$ that maps

complex vertices \mathcal{V} to \mathbb{R}^d ; a combinatorial representation of f as a correspondence map between simplices of the source and target simplicial complexes, M and M'

vertices. We denote by $\mathbf{y}_v, v \in \mathcal{V}$, the column stack of coordinates of v , taken from \mathbf{y} , and by $\text{PL}(M, d)$ we refer to the set of all simplicial maps from $\text{conv}(M)$ to \mathbb{R}^d , defined by (39).

Note that, depending on the context, we use both the combinatorial and geometric representations of simplicial mappings — all these representations are equivalent in our case. For example, an affine mapping f_s of a simplex $s = (v_1, \dots, v_{n+1})$ can be unambiguously defined by specifying images of each vertex v_i under f_s . Generally speaking, we can represent a simplicial map f of a manifold mesh in the following equivalent ways:

1. Simplicial map f can be identified with a list of its affine components $\{f_s | s \in \mathcal{S}\}$, where f_s are affine maps of individual simplices that coincide on common faces of simplices.
2. Simplicial map f can be represented as a functional of the target vertex coordinates \mathbf{x} . That is, $f = f[\mathbf{x}]$, where $\mathbf{x} \in \mathbb{R}^{d|\mathcal{V}| \times 1}$ is a column stack of the vertex coordinates in the target domain and $f[\mathbf{x}]$ is a piecewise affine function that satisfies the following equations⁷:

$$f[\mathbf{x}](\mathbf{y}_j) = \mathbf{x}_j, \quad j = 1, \dots, |\mathcal{V}|. \tag{40}$$

We denote the affine component of $f[\mathbf{x}]$ on simplex $s \in \mathcal{S}$ by

$$f_s[\mathbf{x}] \triangleq (f[\mathbf{x}])_s.$$

3. Simplicial map f can be associated with a function $f_{\mathcal{V}} : \mathcal{V} \rightarrow \mathbb{R}^d$ that maps complex vertices to their target coordinates, i.e., $f_{\mathcal{V}}(v) = \mathbf{x}_v, v \in \mathcal{V}$.
4. Consider the image of f as a (target) simplicial complex $M' = (\mathcal{S}', \mathcal{V}', \mathbf{x})$, i.e., M' is the simplicial complex of

simplices \mathcal{S} , vertices \mathcal{V} and vertex coordinates \mathbf{x} obtained by mapping source coordinates $\mathbf{y}_v, v \in \mathcal{V}$ to \mathbb{R}^d by f . Then, f can be identified with the correspondence map between the source and the target simplicial complexes M and M' . Namely, f can be represented as the mapping between the corresponding source and target vertices, or f can be associated with the mapping between the corresponding source and target simplices.

Fig. 5 illustrates the above equivalent representations of a simplicial map on a triangular mesh, embedded in \mathbb{R}^3 .

3.1 Canonical Representation of Simplicial Maps

Since Jacobian of a local-diffeomorphism $f \in \text{Diff}(\mathbb{R}^n)$ is a non-degenerate matrix, our underlying assumption was that, in the continuous case, source and target domains have the same dimensions.

However, when considering simplicial maps (39), the simplex dimension n and the dimension d of the target domain are not necessary equal.

For this reason, when practical approaches to (1) are considered, one should be aware of certain differences between the problem with equal dimensions, $m = n = d$, and more general scenarios. Following these considerations, we define the source and target *codimensions*

$$\text{codim}(M) \triangleq m - \dim(M); \tag{41}$$

$$\text{codim}(M') \triangleq d - \dim(M'). \tag{42}$$

and, based on these definitions, we consider the following three major scenarios:

$$\text{codim}(M) = 0, \quad \text{codim}(M') = 0; \tag{43}$$

$$\text{codim}(M) > 0, \quad \text{codim}(M') = 0; \tag{44}$$

$$\text{codim}(M) \geq 0, \quad \text{codim}(M') > 0. \tag{45}$$

To treat scenarios (43)-(45) in a unified manner, we consider a local presentation of simplicial map (see Fig. 6). This

⁷ We identify vertices with the indices, $1, \dots, |\mathcal{V}|$, and use square brackets in (40) to indicate that simplicial map f is a function of \mathbf{x} , while round brackets denote the evaluation of simplicial map $f = f[\mathbf{x}]$ at a given point in \mathbb{R}^m .

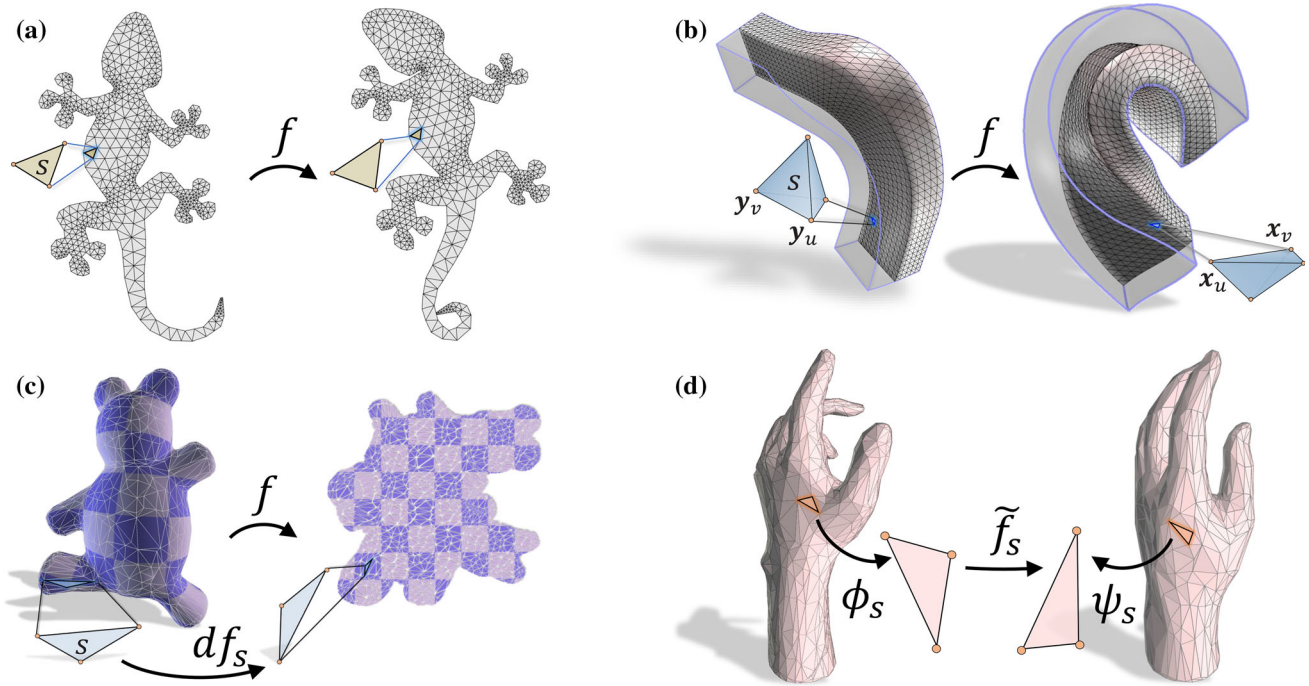


Fig. 6 Showing examples of different types of simplicial mapping and illustrating how distortions under these maps are measured for each simplex. Scenario (43) of mapping between simplicial complexes with zero codimensions in 2D and 3D are depicted in (a) and (b), respec-

tively. An example of scenario (44) is illustrated in (c), where we show a surface flattening and the resulting texture mapping. The most general scenario (45) is illustrated in (d) by showing a map between two triangle meshes, embedded in \mathbb{R}^3 (surface matching problem)

representation removes the extra degrees of freedom, presented in (44) and (45), and it enables a smooth transition between the discrete and continuous settings, presented in Sect. 2. In particular, if codimensions of simplicial complexes in (41) and (42) are non-zero, then we define the distortion energy of a simplicial map $f = f[x]$ by considering “an equivalent map” between n -dimensional simplicial complexes in \mathbb{R}^n . For obtaining an equivalent canonical map, we project f ’s components onto the n -dimensional subspace, without distorting shapes of the source and target simplices.

Assume that f is a simplicial map, defined in (39), and that $n \leq d, n \leq m$. Denote by s' the target simplex of $s \in \mathcal{S}$ under f and let ϕ_s and ψ_s be rigid transformations of \mathbb{R}^m and \mathbb{R}^n that map $\text{conv}(s)$ and $\text{conv}(s')$ into hyperplanes $\mathbb{R}^n \times \{0\}^{m-n}$ and $\mathbb{R}^n \times \{0\}^{d-n}$, respectively. To simplify our notations, we assume w.l.o.g. that these hyperplanes are equal to \mathbb{R}^n . If \mathcal{D} is a first order distortion of deformations $\text{Diff}(\mathbb{R}^n)$, then, we define

$$\mathcal{D}(f_s) \triangleq \mathcal{D}([\tilde{f}_s]_{\sim}, \mathbf{r}), \quad s \in \mathcal{S}, \mathbf{r} \in \text{conv}(s),$$

where $[\cdot]_{\sim}$ is an equivalence class of deformations in \mathbb{R}^n , defined by (19), and \tilde{f}_s is the transition map

$$\tilde{f}_s \triangleq \psi_s \circ f_s \circ \phi_s^{-1}. \tag{46}$$

As illustrated in Fig. 6d, map \tilde{f}_s satisfies the following diagram:

$$\begin{array}{ccc} \text{conv}(s) \subset \mathbb{R}^m & \xrightarrow{f_s} & \text{conv}(s') \subset \mathbb{R}^d \\ \downarrow \phi_s & & \downarrow \psi_s \\ \mathbb{R}^n & \xrightarrow{\tilde{f}_s} & \mathbb{R}^n \end{array} .$$

We refer to (46) as to the (local) canonical representation of a simplicial map $f = f[x]$. We use canonical representation of f to define distortion energy of \mathcal{D} ,

$$E_{\mathcal{D}}(f[x]) \triangleq \sum_{s \in \mathcal{S}} w(s) \mathcal{D}([\tilde{f}_s]_{\sim}), \tag{47}$$

where $w(s)$ are simplex weights, most frequently chosen to be the simplex volume $\text{Vol}(s)$. We always compute distortions of canonical representations of simplicial maps. Obviously, substituting components of a map $f \in \text{PL}(\mathbf{M}, d)$ with its canonical representations does not change distortion values.

Noteworthy is the fact that the energy of (47) can be equivalently represented by other expressions, e.g., using a vertex-based weighted sum of distortion densities [24,81]. However, vertex weights and other related representations are rarely used in practical applications. In this paper, we




Fig. 7 An example of GD optimization of isometric distortion induced by mapping of a 2-dimensional simplicial complex from \mathbb{R}^2 (left) to \mathbb{R}^3 (right). We initialize the problem by mapping the source mesh onto

a sphere (target) and execute the GD optimization without using positional constraints on vertex coordinates

follow a more common definition (47) which is based on the linear finite element formulation. By using these formulations in the discrete case, we can employ a more general optimization framework.

4 Minimizing Distortions of Simplicial Maps

Although we have provided a unified method for computing distortion energies in scenarios (43)-(45), these scenarios require a different treatment for minimizing distortions.

To illustrate fundamental differences between distortion optimization techniques, employed in scenarios (43) and (45), consider a gradient descent (GD) optimization of a distortion energy (47). If $\text{codim}(M') = 0$, then the GD update of target vertices, $\mathbf{x} \leftarrow \mathbf{x} - \Delta t \nabla_{\mathbf{x}} E$, is well defined, since the displacement vector $\Delta \mathbf{x}$ and the target simplicial complex M' belong to the same n -dimension plane, $\mathbb{R}^n = \mathbb{R}^d$.

In contrast to the case $\text{codim}(M') = 0$, if $\text{codim}(M') > 0$, then there is a certain ambiguity in the GD optimization of \mathbf{x} , since it should be specified whether the target vertex coordinates $\mathbf{x}_j, j = 1, \dots, |\mathcal{V}|$, are free to move in \mathbb{R}^d (free-form deformations), or \mathbf{x}_j are constrained to be contained in a given n -dimensional submanifold $T^n \subset \mathbb{R}^d$.

In practice, scenario (45) with free-form deformations is rarely processed by optimizing directly distortions on M . This is because distortion energy (47) by itself cannot control the shape of a target domain, since, in this case, there are extra degrees of freedom to rotate separately simplex images in \mathbb{R}^d without changing distortions. See Fig. 7 for an example of those adverse effects that appear in a free-form deformation of a triangular mesh.

For instance, computing optimal mapping between two triangular meshes, embedded in \mathbb{R}^3 , is a problem of type (45) in which \mathbf{x} is constrained to be contained in a given 2-manifold. This problem has various applications in computer

vision and imaging. In particular, the existing approaches to that problem are employed for such tasks as shape matching and surface registration. Often, shape matching algorithms combine both distortion minimization methods [6,31,93] and other geometry processing techniques such as: computations of geodesic distances [97], functional maps [30,84] and etc.

If $\text{codim}(M') > 0$, then a practical approach to minimizing distortions on M , with unconstrained coordinates $\mathbf{x} \in \mathbb{R}^{d|\mathcal{V}| \times 1}$, is to consider simplicial maps of d -dimensional complex D containing the original n -dimensional complex M . Then, a solution of problem (1) on D is an optimal simplicial map

$$f_D^* : \text{conv}(D) \rightarrow \mathbb{R}^d,$$

of type (43) or (44), and the restriction of f_D^* to $\text{conv}(M)$ induces a low distortion simplicial map f_M^* of the complex M (see the illustration in Fig. 8). In particular, if $v \in \mathcal{V}$ is contained in a d -dimensional simplex c of D , then the image $f_M^*(v)$ can be computed by using barycentric coordinates of v in c .

Next, we discuss how to compute Jacobian matrices of simplicial maps by using barycentric coordinates and other related quantities. Note that, unless stated otherwise, we further assume that $\text{codim}(M') = 0$, i.e., $d = n$.

4.1 Jacobian Computation

On the one hand, the problem of minimizing (47), in all practical situations, is stated in terms of vector \mathbf{x} and a simplicial complex M . On the other hand, the energy, through distortion density, depends on the Jacobians $df_s, s \in \mathcal{S}$. Therefore, from a practical viewpoint, it is important to develop concise analytical expression for

$$df_s : \mathbb{R}^{n|\mathcal{V}|} \rightarrow \mathbb{R}^{n \times n},$$




Fig. 8 An optimal mapping of a volumetric domain D (left) that induces a low distortion mapping of a surface mesh M (right), contained inside the volume. The volumetric mapping f_D was computed by minimizing isometric distortion over a tetrahedral mesh. The triangulated surface M was represented as a volumetric texture inside D . Then, vertices of M were mapped by f_D according to their barycentric coordinates

in terms of the known quantities⁸ \mathbf{x} and \mathcal{S} . First, we consider one-dimensional maps $\text{PL}(M, 1)$. This space has a natural basis of Lagrange basis functions $\{h_v \in \text{PL}(M, 1) : v \in \mathcal{V}\}$, also called *hat functions*, where each h_v satisfies the following system

$$h_v(\mathbf{y}_u) = \delta_{vu}, \quad v, u \in \mathcal{V}, \tag{48}$$

$$\sum_{v \in \mathcal{S}} h_v(\mathbf{r}) = 1, \quad s \in \mathcal{S}, \mathbf{r} \in \text{conv}(s). \tag{49}$$

Representing simplicial map $f[\mathbf{x}] \in \text{PL}(M, n)$ with respect to basis functions h_v yields

$$f_s(\mathbf{r}) = \sum_{j=1}^{n+1} h_{v_j}(\mathbf{r}) \mathbf{x}_{v_j}, \quad s \in \mathcal{S}, \mathbf{r} \in \text{conv}(s). \tag{50}$$

The close form solution of (49) and (48) has a simple geometric interpretation in \mathbb{R}^2 and \mathbb{R}^3 . Let $s = (v_1, \dots, v_{n+1}) \in \mathcal{S}$, $\mathbf{r} \in \text{conv}(s)$ and denote by μ_j the face of s located opposite to vertex v_j , and let $\boldsymbol{\eta}_j$ be a vector normal to μ_j whose length equals $\|\boldsymbol{\eta}_j\| = \text{Area}(\mu_j)$ (see Figs. 9 and 10). Then, we can express n basis functions as

$$h_{v_j}(\mathbf{r}) = \frac{(\mathbf{r} - \mathbf{y}_{v_{n+1}}) \cdot \boldsymbol{\eta}_j}{n \text{Vol}(s)}, \quad j = 1, \dots, n, \tag{51}$$

and, for attaining a convex combination in (49), the last basis function is set to

$$h_{v_{n+1}}(\mathbf{r}) \triangleq 1 - (h_{v_1}(\mathbf{r}) + \dots + h_{v_n}(\mathbf{r})).$$

In fact, $(h_{v_1}(\mathbf{r}), \dots, h_{v_{n+1}}(\mathbf{r}))$ are *barycentric coordinates* of point \mathbf{r} in s . These coordinates are widely employed in geometric computer vision for interpolating discrete quantities, sampled at vertices. By differentiating (50) with respect

⁸ Here we use the shorthand notation $\mathbb{R}^{n|\mathcal{V}|} = \mathbb{R}^{n|\mathcal{V}| \times 1}$.




Fig. 9 Illustration of a hat function (51), defined over a 2D simplicial complex, embedded in \mathbb{R}^3

to \mathbf{r} , we obtain the Jacobian df_s

$$df_s = \sum_{j=1}^{n+1} \frac{\partial h_{v_j}}{\partial \mathbf{r}} \mathbf{x}_{v_j}. \tag{52}$$

Gradients $\partial h_v / \partial \mathbf{r}$ are constant for each $v \in s \in \mathcal{S}$ because $h_v|_{\text{conv}(s)}$ are affine functions, by the definition. Hence, df_s is constant in $\mathbf{r} \in \text{conv}(s)$ and the correspondence between the target coordinates and Jacobian of f_s forms a linear operator $\mathbb{R}^{n|\mathcal{V}|} \rightarrow \mathbb{R}^{n \times n}$, denoted by

$$\partial_s(\mathbf{x}) \triangleq df_s[\mathbf{x}], \quad s \in \mathcal{S}.$$

The details of these computations are presented in Appendix A.

4.2 Problem Formulation

We are now in the position to provide a formal definition of the discrete problem, based on definitions of simplicial maps and evaluations of their Jacobians. Let $M = (\mathcal{S}, \mathcal{V}, \mathbf{y})$ be a simplicial complex, we then consider the following problem of optimizing distortions in the discrete settings:

$$f^* = \underset{f \in \text{PL}(M, n)}{\text{argmin}} E(f); \tag{53}$$

$$\text{s. t. } \det(df_s) > 0, \quad \forall s \in \mathcal{S}; \tag{54}$$

$$f_{\mathcal{V}}(a_i) = \mathbf{x}_{a_i}^0, \quad i = 1, \dots, k, \tag{55}$$

where (53) is the optimization problem of distortion energy, defined according to (47); (54) is the orientation constraints on maps $\text{PL}(M, n)$ and (55) are discrete positional constraints defined on a set of vertices (*anchor points*) a_1, \dots, a_k which coordinates are fixed to given positions $\mathbf{x}_{a_1}^0, \dots, \mathbf{x}_{a_k}^0$.

It is more convenient to represent simplicial maps by (40) and to formulate constraints via a system of linear equations because it reduces (53), defined over $\text{PL}(M, n)$, to a more simple optimization problem in $\mathbb{R}^{n|\mathcal{V}|}$:

$$\mathbf{x}^* = \underset{\mathbf{x} \in \mathbb{R}^{n|\mathcal{V}|}}{\text{argmin}} E(f[\mathbf{x}]); \tag{56}$$

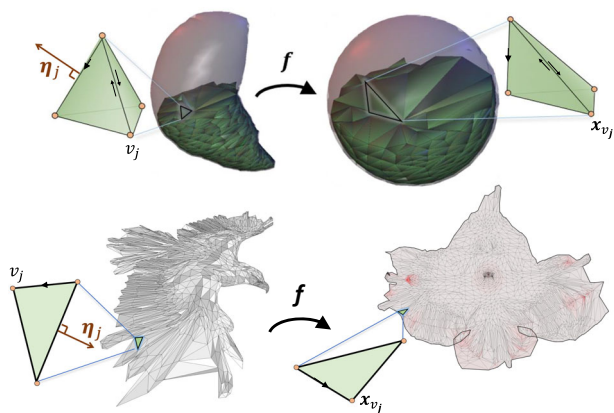


Fig. 10 Examples of simplicial mapping from \mathbb{R}^3 to \mathbb{R}^2 (bottom) and from \mathbb{R}^3 to \mathbb{R}^3 (top). The latter example demonstrates a volumetric parameterization of the segment of a CT scan (hippocampal region) [78]. Highlighted simplices illustrate a piecewise affine construction, introduced in (51)

$$\text{s. t. } \det(df_s[\mathbf{x}]) > 0, \forall s \in \mathcal{S}; \tag{57}$$

$$A\mathbf{x} = \mathbf{b}, A \in \mathbb{R}^{k \times n|\mathcal{V}|}, \tag{58}$$

where (58) is the generalization of (55) to linear positional constraints and the objective energy is the real function $E = E(\mathbf{x}) : \mathbb{R}^{n|\mathcal{V}|} \mapsto \mathbb{R}$.

We proceed to review various methods for solving the distortion energy equation (56).

4.3 Distortion Optimization Methods

In practice, solutions of (56) are orientation-preserving maps of triangular and tetrahedral meshes, embedded in \mathbb{R}^2 and \mathbb{R}^3 . Various iterative algorithms and other related techniques are used for constructing optimal deformations of triangular and tetrahedral meshes via minimization of geometric energies.

The relevant methods can be qualitatively divided into the four major categories: (i) linear methods which, in general, are the earliest and fastest methods; (ii) convexification methods; (iii) non-linear optimization techniques; (iv) indirect approaches, intended for a limited subset of objective energies with a well studied structure.

Linear methods. These methods compute simplicial maps by solving a linear system

$$L\mathbf{x} = \mathbf{b}, \tag{59}$$

where the coordinates \mathbf{x}_v for each vertex $v \in \mathcal{V}$ are expressed by a weighted average of its neighbors. The matrix L is often considered as a discrete approximation of the Laplace-Beltrami operator. Consequently, solutions of (59), called *discrete harmonic maps*, are minimizers of a piecewise Dirichlet energy.

Among others, the most common weighting schemes, employed in (59) for triangular meshes, are uniform, cotangent weights and weights derived from the mean-value coordinates [34]. Whenever these weights satisfy $L_{uv} > 0$ for any neighboring vertices $u \neq v$, and coordinates of boundary vertices in \mathbf{b} form a convex polygon, then solving (59) yields an injective mapping $f[\mathbf{x}]$ of a source mesh into the plane [33].⁹ Furthermore, certain methods for injective harmonic mapping can be extended from target planar domains to more general domains on surfaces [2,4,5].

If L is a cotangent weighted Laplacian matrix and \mathbf{x} is the solution of (59) with respect to L , then $\mathbf{x}^\top L\mathbf{x}$ is the discrete Dirichlet energy induced by (25). Moreover, according to [34], a simplicial map $f[\mathbf{x}], \mathbf{x} = L^{-1}\mathbf{b}$, attains conformality if the density of the energy of $\mathcal{D}_{\text{Dirichlet}}$, per a unit area, reaches its global minimum.

Levy et al. [72] simplify the non-linear MIPS optimization by introducing a linear approach for computing least-square conformal maps. More recent study of [10] unifies major least squares geometry processing techniques into a framework of shape projection operators. The recent study of [20] computes conformal parameterization by partitioning surface into smaller subdomains and computing linear conformal flattening of each subdomain.

Although shape operators and the related linear techniques have a low computational cost, these methods are limited in their application to a narrow set of geometric measures and support only certain positional constraints. This is in contrast to recent studies in non-linear optimization and convexification methods, aimed at minimizing arbitrary distortions expressed by Jacobian singular values.

Convexification methods. In general, these are iterative techniques that approximate problems similar to (53) by a sequence of nested convex problems, for which convex optimization tools can be applied. These include: projecting simplicial mappings onto the space of Bounded Distortion (BD) maps of triangular [66] and tetrahedral [3] meshes; Large Scale Bounded Distortion (LBD) maps [56]; controlling singular values via Semi-Definite Programming (SDP) [55].

The space of K -bounded distortion mappings is defined as the set of simplicial maps $f \in \text{PL}(\mathcal{M}, n)$ satisfying the inversion-free constraints (54) and such that the conformal distortion (29) of linear components of each f is bounded by a given number $K > 1$,

$$\mathcal{D}_{\text{conf}}(df_s) \leq K, \forall s \in \mathcal{S}.$$

The method of [66] for bounded distortion mapping of triangular meshes is based on

⁹ See [33] for additional constraints on the mesh connectivity.

representation of affine transformations using complex numbers and on the following formula for singular values of 2×2 Jacobian matrices df_s :

$$\sigma_1(df_s) = |\alpha_s| + |\beta_s|; \tag{60}$$

$$\sigma_2(df_s) = \left| |\alpha_s| - |\beta_s| \right|, \tag{61}$$

where $\alpha_s, \beta_s \in \mathbb{C}$ are complex representations of the similarity and anti-similarity components of $df_s, s \in \mathcal{S}$. As shown in [66], the maximal convex subset of K -bounded distortion mappings in 2D can be characterized by a system of simple inequalities, expressed in terms of $K, |\alpha_s|$ and $|\beta_s|$, for $s \in \mathcal{S}$.

The subsequent work of [3] extends this strategy to tetrahedral meshes, by solving a quadratic problem of projecting a given simplicial map onto the BD space. This projection technique can be further improved by employing KKT linear systems for a more efficient formulation of optimization constraints [56].

Convexification methods can incorporate non-locally injective initializations. In particular, the BD and LBD methods are aimed primary at repairing non-positively oriented maps and restrain their distortions within a finite range, whereas SDP, and other related interior point solvers, are considered to be not sensitive to the quality of an initial map.

Similarly to other optimization techniques, convexification methods express objective measures in terms of the Jacobian singular values. However, unlike more general optimization algorithms, such as gradient descent or quasi-Newton solvers, convex optimization tools, employed in convexification methods, impose much tougher constraints on the optimization process.

Next, we discuss a more generic optimization approaches, referred as to methods in *non-linear geometric optimization*. While convexification solvers employ convex approximation both for the objective function $E(f)$ and for the space of objective variables $f \in \text{PL}(\mathcal{M}, n)$, non-linear methods approximate energy (47), alone, and thereby these methods can be applied in more general scenarios.

Non-linear geometric optimization. A typical non-linear geometric solver updates the mapping $f[\mathbf{x}^i]$ at the i^{th} iteration as follows

$$\mathbf{x}^{i+1} = \mathbf{x}^i + \Delta t^i \mathbf{d}^i (\nabla_{\mathbf{x}} E, H_{\mathbf{x}^i}), \tag{62}$$

where \mathbf{d}^i is the field of the descending direction computed as a function of the distortion gradient $\nabla_{\mathbf{x}} E = \nabla_{\mathbf{x}} E(\mathbf{x}^i)$ and the Hessian $H_{\mathbf{x}^i}$. Here we considered iterative solvers that

begin with an initialization $f^0 = f[\mathbf{x}^0]$ and recompute (62) for each iteration¹⁰ i .

The exact amount of Δt^i by which \mathbf{x}^i is modified along the descent fields is computed, in general, by inexact line search; e.g., using the Armijo back-tracking method.

Based on methods of computing descent direction, geometric solvers are divided into the *first* and *second* order techniques.

The gradient descent (GD) is the basic first order method in which vertex coordinates are updated along the negative gradient direction, i.e.,

$$\mathbf{d} = -\nabla_{\mathbf{x}} E.$$

It is easy to implement GD for any SVD-based distortion \mathcal{D} because the only steps needed for GD update are: the computation of $\nabla_{\mathbf{x}} E$ and the line search step. Moreover, for each $v \in \mathcal{V}$ the gradient component $(\nabla_{\mathbf{x}} E)_v$ depends only on the neighboring vertices of v . [See details on gradient computation in Appendix A.2.] Therefore, it is easy to apply GD, locally, to minimize energy (47) over a small subset of simplices. Local GD methods, in which only a fraction of vertices are updated at each iteration, are called *block gradient descent* (BGD). Due to a simple implementation and robustness, GD and BGD algorithms are widely employed in geometry processing; these methods were popular in early applications [45] and they also are used in recent studies for inducing locally injective simplicial mappings [37,81,83]. Although GD works well over small blocks of vertices, it converges slowly when applied to the coordinates of all vertices in (62).

To speed up distortion optimization, one should employ more general first order solvers in which \mathbf{d} is the solution of a sparse linear system

$$\tilde{H}_{\mathbf{x}^i} \mathbf{d} = -\nabla_{\mathbf{x}} E(\mathbf{x}^i), \tag{63}$$

where \tilde{H} is the Hessian of a quadratic approximation of the original energy E in the vicinity of \mathbf{x}^i (*quadratic proxy*):

$$\begin{aligned} \tilde{E}(\mathbf{x}) &= E(\mathbf{x}^i) + (\mathbf{x} - \mathbf{x}^i)^T \nabla_{\mathbf{x}} E(\mathbf{x}^i) \\ &\quad + \frac{1}{2} (\mathbf{x} - \mathbf{x}^i)^T \tilde{H}_{\mathbf{x}^i} (\mathbf{x} - \mathbf{x}^i). \end{aligned} \tag{64}$$

For example, in Sobolev gradient descent (SGD), and in the closely-related Accelerated Quadratic Proxy (AQP) [57], \tilde{H} is a cotangent-weighted mesh Laplacian L , computed over source coordinates \mathbf{y} . These methods are designed for minimizing isometric distortions. Further, the method of

¹⁰ It is often clear, from the context, at what iteration vertex coordinates and descent fields are computed. Thereby, to make our presentation more simple, we often drop the superscript indices, used in (62).




Fig. 11 Isometric parameterization of a Hilbert fitting curve (source mesh) with different optimization methods. The surface is flattened by minimizing symmetric Dirichlet energy (32), initializing the process by

the Tutte embedding of the source mesh into a disc. We depict parameterization results, obtained at different iterations, and show the final texture mapping at the right side

Scalable Locally Injective Mappings (SLIM) [89] extends SGD approach to general distortions by setting \tilde{H} to be a reweighted Laplacian. In the isometry-aware preconditioning method [17] (AKVF), \tilde{H} is a quadratic form of the Killing energy of vector fields, defined over triangle meshes.

Similarly, second order methods compute \tilde{H} as a function of both ∇E and $\nabla^2 E$. Note that \tilde{H} should be a positive semidefinite matrix to guarantee that \mathbf{d} , computed in (63), is a descent direction of $E(\mathbf{x})$. Therefore, second order methods cannot employ the original Hessian $\tilde{H} = \nabla^2 E$ because, in most cases, energy E is non-convex in \mathbf{x} (see Sect. 8). Instead, second order methods use second order derivatives of $E(\mathbf{x})$ to approximate the Hessian. For example, $\tilde{H} = \text{diag}(\nabla_{\mathbf{x}}^2 E)$ in Jacobi gradient descent [117]; $\nabla^2 E$ is projected into a positive semidefinite cone [63, 69, 108] in projected Newton methods (PN). In 2D, a positive semidefinite approximation of $\nabla^2 E$ can be computed by using complex analysis [26, 42] or via the Composite Majorization (CM) method [102], in which \tilde{H} is derived from the analytic expression of singular values of matrices in $\mathbb{R}^{2 \times 2}$ (60) and (61).

Quasi-Newton methods lay in between first and second order solvers. Similarly to second order solvers, quasi-Newton methods are based on a Newton update step. However, instead of directly computing second order derivatives, quasi-Newton methods use gradients and vertex positions from previous iterations to iteratively update approximate Hessians, $\tilde{H} = \tilde{H}(\nabla E^i, \nabla E^{i-1}, \dots, \mathbf{x}^i, \mathbf{x}^{i-1}, \dots)$. Due to their robustness and relatively low computational cost, these methods are widely employed in geometry processing and

imaging. For example, Smith et al. [103] had adopted the classical L-BFGS algorithm for computing globally bijective parameterization. This method is enhanced in the sequel by the Blended Cured Quasi-Newton (BCQN) [121] strategy of a gradual blending between AQP [57] and L-BFGS. BCQN solver benefits both from the super-linear convergence of L-BFGS in the vicinity of a local minimum and from the rapid progress of AQP at the first iterations.

In Fig. 11, we demonstrate a number of non-linear geometric solvers, employed for parameterization of a triangulated surface. As shown by the figure, iterations of second order methods result in a more rapid progress. However, computing descent direction is more costly in second order solvers, since the Hessian approximation in these methods involves computations of both the first and second order energy derivatives.

Indirect approaches. So far, distortion optimization methods, presented in this section, have the following common property: these approaches compute optimal simplicial maps in which the measure of map optimality is a function of the vertex coordinates \mathbf{x} . However, in many applications, simplicial maps are represented implicitly. For example, simplicial maps can be represented implicitly as a solution of some equation, as a realization of a discrete metric and etc. In these cases, the final optimal map is reconstructed from its implicit presentation and the measure of the map's optimality is a property of the implicit form, used to define the map. We refer to methods, based on implicit representations, as to *indirect approaches*.

Although presenting all indirect approaches is beyond the scope of this paper, we list some indirect methods that are closely related to the distortion minimization problem. Some of these methods are discussed in more details in Sect. 7.

For example, discrete quasi-conformal maps can be constructed on triangulated surfaces by solving equations of Beltrami coefficients [19,122]. A discrete approximation of conformal maps can be computed in a metric domain, e.g., discrete Ricci flow obtains a user-specified distribution of Gaussian curvatures at vertices and it achieves the targeted metric conformal to the original metric. Some indirect approaches represent triangulated surfaces by curvature values or by unit normals, defined on vertices. These methods compute implicit representations of optimal maps by modifying these values on vertices, e.g., by performing the conformal curvature flow [23], or by computing the unit normal flow [125].

Another important class of indirect approaches presents discrete methods for computing Optimal Mass Transportation (OMT). In general, methods in OMT seek to find a volume/area-preserving map between two spaces that minimizes a specific transportation cost. There are multiple equivalent ways to define such problems on meshes, including the classical Monge-Kantorovich formulation, the Wasserstein metric formulation and the Brenier formulation for OMT in the semi-discrete case (see the survey of [101]). Although transportation cost functions can be formulated in a different manner than distortion measures, there exist many interrelated methods for OMT and distortion minimization. For example, [124] introduced an algorithm for area-preserving flattening of triangle meshes. First, the algorithm is initialized with conformal map f^0 , it then minimizes area distortions by solving an OMT problem in which conformal factors of f^0 define the transportation cost. Likewise, OMT and conformal mappings are used in many shape analysis applications for comparing multiple objects and detecting geometric change in an object. For instance, the dissimilarity measure of two meshes M and M' with disc topology (or two genus-zero meshes) can be defined as the minimal transportation cost induced over mappings $\{(f')^{-1} \circ g \circ f\}$, where f and f' are conformal parameterizations of M and M' into a disc (sphere), and g is a Möbius transformation [12,71].

Worth mentioning, minimizing variants of volume distortions yields a more general class of *density-equalizing maps*. Instead of preserving the volume, density-equalizing maps preserve a given volume density. Such variants of volume-preserving maps are used in classical tasks of data visualization [28,41]. For instance, Choi and Rycroft [24] introduced a method for density-equalizing mapping of triangular meshes. The method operates by flattening simply-connected meshes in a way that inflates or shrinks the target mesh triangles according to the specified area densities. This

method starts with curvature-based flattening of the mesh boundary, then it uses a diffusion-based algorithm to minimize area densities over vertices.

4.4 Acceleration Techniques

There is a number of techniques aimed at accelerating existing algorithms for distortion minimization. For example, Accelerated Quadratic Proxy [57] (AQP) employs a Nesterov-like acceleration to boost the vertex update step (62). First, AQP computes x^{i+1} according to (62), and then it sets x to be an affine combination of x^{i+1} and the target coordinates obtained at the previous iteration. That is, $x = (1 + t)x^{i+1} - tx^i$, for a small $t \in (0, 1)$.

Instead of using the source mesh as a reference, Progressive Parametrization (PP) [74] defines an intermediate (progressive) mesh that induces low isometric distortion and that is as-close-as-possible to the given source domain. PP accelerates non-linear solvers by iteratively generating progressive meshes and by decomposing source-to-target map into intermediate mappings with bounded singular values.

Anderson Acceleration (AA) [85,86,123] is another approach to speeding up optimization, based on a well-established technique for fixed-point solvers. Methods of AA are designed for alternating local-global optimization, such as [10,13,73]. A local-global algorithm alternates between the two steps: (i) the local step, in which energy E is minimized with respect to auxiliary variables, while the target coordinates x are fixed; (ii) the global step, in which E is minimized with respect to x . The key idea of AA is to speed up the global step by modifying auxiliary variables, obtained at the local step. For example, AA can be used to improve existing methods for BD mapping [98].

Another class of popular acceleration techniques is based on hierarchical mappings between meshes of decreasing resolution. For example, ABF++ algorithm [100] of conformal flattening performs a sequence of edge collapse operations for decimating the original mesh M . This method computes a low resolution parameterization f' of decimated mesh M' , and then it derives f' for the original mesh by applying a series of vertex split operations. In certain scenarios, optimization (56) can be initialized directly with a solution $(f')^*$ of the same problem, computed in a low resolution. As shown in [82,83], $(f')^*$ can be transformed to an inversion-free initialization of M by, first, mapping M into a disc, and then deforming that disc into the shape of $\text{Img}((f')^*)$.

5 Distortion Optimization Constraints

This section is dedicated to a detailed analysis of the optimization constraints (58) and (57). We also discuss existing methods for locally and globally injective mapping, and

relate these methods to orientation constraints and to distortion minimization under fixed boundary constraints.

5.1 Positional Constraints

Positional constraints (58) constitutes an integral part of the linear methods. Moreover, computing unconstrained harmonic maps, in most linear methods, such as [109] or [38], yields a trivial solution that contracts the entire target domain into a single point.

In contrast to linear techniques, convexification and non-linear optimization methods can be used without positional constraints. For example, surface parameterization with a free boundary can be computed as a solution of (56) without positional constraints (Figs. 11 and 2 (right)).

Positional constraints (58) can be integrated into first and second order non-linear solvers via the constrained minimization of the quadratic proxy (64). In particular, if \tilde{H}_x is the Hessian approximation at x , then the constrained descent direction is computed via the following KKT system:

$$\begin{pmatrix} \tilde{H}_x & A^\top \\ A & 0 \end{pmatrix} \begin{pmatrix} d \\ \lambda \end{pmatrix} = \begin{pmatrix} -\nabla E_x \\ b \end{pmatrix}, \tag{65}$$

where (A, b) are linear constraints (58) and λ is a KKT multiplier vector.

It is very common to solve (53) with *Dirichlet constraints*, that is, with constraints (58), reduced to a set of fixed anchors, $x_{a_i} = x_{a_i}^0, a_i \in \mathcal{V}_{\text{fixed}} \subset \mathcal{V}$ (see Fig. 18). In this scenario, system (65) is reduced to solving descent direction equation (63) over the set of free vertices $\mathcal{V} \setminus \mathcal{V}_{\text{fixed}}$. Particularly, in gradient based methods, such as GD and BGD, the Dirichlet constraints can be implemented by modifying only free vertices, while constrained vertices are fixed at their prescribed positions.

5.2 Orientation Constraints

Unlike the positional constraints, orientation constraints are non-linear and thus cannot be processed directly by (65). For satisfying (57), one should be able to either repair inverted simplices, or to start with an inversion-free map and preserve its positive orientation along the entire optimization.

In general, the process of repairing an inverted simplex s by (62), has to collapse s before unfolding it to a positive orientation in the target domain. However, most popular distortion energies are regular distortions that satisfy the barrier properties from Definition 2.12. These distortions explode on degenerate target simplices,¹¹ and thus, optimization of these distortions prevents orientation flips, keeping elements

at their initial orientation. For a non-regular distortion \mathcal{D} , orientation flips can be prevented by adding a barrier term (34) to \mathcal{D} .

As a result, the vast majority of existing geometric optimization algorithms, process orientation constraints, implicitly, by requiring, first, a positively-oriented initialization $f^0 = f[x^0]$ and, then, preserving initial simplex orientation at each subsequent iteration.

The strategies, proposed to keep f satisfying (57), include: designing distortions with flip barriers [37,103] (i.e., using regular distortion or adding a barrier term), inversion-aware line search [103], the barrier-aware line search filtering [121], employment of scaffold meshes [53,69] and their variants [107]. In addition to these methods, one can employ ray-tracing algorithms [81,103] to prevent target simplices from being inverted during the line search stage. In certain cases, an inversion-free initialization x^0 of (56) can be computed by linear methods. For example, unconstrained parameterization of triangle meshes with disc-topology can be initialized by mapping the source mesh into convex planar domain via the classical Tutte embedding [109] or by using linear conformal maps [38]. Although these maps attain a positive orientation, $\det df_s^0 > 0$, and low conformal distortions, they produce highly distorted elements with respect to other measures, such as isometric distortions. Computing feasible initialization f^0 in more general scenarios is much more challenging. To the best of our knowledge, there is no robust solution of providing an inversion-free initialization for problem (56) in a general scenario.

In a view of the above limitations, a number of strategies have been proposed for generalizing existing methods of injective mappings. In particular, a number of approaches have been proposed to extend Tutte’s embedding to non-convex domains.

The method of [119] starts with uniformly weighted Laplacian and iteratively modifies it using cotangent weights, computed with respect to the solution of (59) on the previous step. This approach can be considered as hybrid linear algorithm, aimed at repairing inverted triangles by means of minimizing the deviation between the target unsigned area

$$\text{Area} |f| \triangleq \sum_{s \in \mathcal{S}} |\text{Area}(f_s(s))|$$

and target signed area

$$\text{Area}(f) \triangleq \sum_{s \in \mathcal{S}} \text{Area}(f_s(s)),$$

where $\text{Area}(f_s(s))$ denotes the signed area of the image of simplex s under the affine map f_s .


Assume that we seek to find an orientation-preserving map under a properly set boundary constraints, i.e., under

¹¹ A n -dimensional simplex is degenerate if its n -volume is zero.

the Dirichlet constraints (55), $f_V^0(a_i) = \mathbf{x}_{a_i}^0$, where a_i are indices of boundary vertices and f_V^0 is an inversion-free map. In such a case, the method of [119] can be simplified to yield the following minimization

$$\operatorname{argmin}_{\mathbf{x}} \operatorname{Area} |f[\mathbf{x}]|. \quad (66)$$

Problem (66) is applicable both to triangular and tetrahedral meshes. Clearly, in the latter case, $\operatorname{Area} |f[\mathbf{x}]|$ denotes the total unsigned volume of target tetrahedrons under f . If \mathbf{x}^* is a global minimizer of (66), then $\operatorname{Area} |f^*|$, induced by $f^* = f[\mathbf{x}^*]$, equals to the area (volume) contained in the mesh boundary $\partial \mathbf{x}^*$. In such a case, $\det df_s[\mathbf{x}^*] \geq 0$ for all $s \in \mathcal{S}$, since, as illustrated by the inset, any inverted simplex s with non-zero volume, either intersects its positively oriented neighbors (inner flips), or the target shape of s goes beyond the mesh boundary (outer flips). In both scenarios, $\operatorname{Area} |f^*|$ exceeds $\operatorname{Area} (\partial \mathbf{x}^*)$, contradicting the global optimality of \mathbf{x}^* . However, a simplicial map f^* can produce collapsed simplices and still be a global minimizer of the total unsigned volume. This limitation and the related issue of vanishing gradients $\nabla_{\mathbf{x}} \operatorname{Area} |f|$ are overcome in the recent work of [27] by lifting simplices into the higher dimensional space \mathbb{R}^{2n} . The sum of unsigned areas of lifted simplices is called the Total Lifted Content (TLC). As proved by [27], minimizers of TLC satisfy constraints (57) and thus induce injective mappings of triangular and tetrahedral meshes into domains with non-overlapping boundaries (see Sect. 5.3). Nevertheless, similar to algorithms in harmonic mappings, TLC requires a proper fixation of all boundary vertices. Minimizing TLC with more general positional constraints, such as fixations of a small number of anchor points (Fig. 6a), can cause inversions and undesirable shrinking of mesh elements. Hence, in such tasks as shape deformation and parameterization with non-fixed boundaries, one should use other approaches to computing inversion-free initializations. Those include a number of linear and non-linear methods, presented below.



Sawhney and Crane [94] have proposed a boundary First Flattening (BFF) as a linear method for conformal flattening of triangular meshes that supports direct manipulation of lengths or angles of boundary edges. Although theoretically BFF is guaranteed to produce inversion-free flattening only onto convex domains, in practice it is capable of producing positively oriented maps onto simple non-convex domains.

If the algorithm, used to initialize (56), fails to produce an inversion-free map, then the existing inverted simplices can be repaired via a limited number of convexification methods, such as BD [3,66] and LBD [56]. These methods project a given map f onto the nearest positively-oriented simplicial map with bounded conformal distortions. However, there are no guarantees to keep constraints (55) under these projections. Setting a suitable lower bound K of conformal distortion is another common drawback of algorithms in BD mapping. For instance, the K -bounded distortion space can be empty if the value of K is too low. Likewise, setting K too high may increase significantly the number of required projections. As shown in [98], some of the above issues of BD mapping can be resolved by an iterative modification of the bound K and by employing local-global acceleration techniques for geometric optimization [86].

In certain cases, optimization of non-barrier distortion measures can be initialized with orientation-reversing maps. For instance, Weber et al. [115] proposed a non-linear minimization of the Least Square Beltrami (LSB) energy for inducing extremal quasi-conformal mapping of triangle meshes. This algorithm can be initialized with foldovers because LSB energy is finite on collapsed and inverted triangles. Since this method is based on the Teichmüller spaces of conformal equivalence classes of surfaces [54], it cannot be extended from triangular to tetrahedral meshes. For the same reason, there is no good analogy of LSB energy for non-conformal distortion measures.

If \mathcal{D} is a general barrier-type distortion, then one of the practical ways to allow orientation-reversing initialization of (56) is to remove barrier terms of \mathcal{D} on inverted simplices [83]. Fig. 12 shows an example of how this modification impacts the isometric distortion (32).

Simplex assembling is another approach for enforcing a consistent orientation in problem initialized with non-locally injective maps. In simplex assembling methods [36,87], meshes are disassembled into topologically disconnected subsets. Self-intersections of simplices are repaired, first, by minimizing inversion penalties in each component. Then, the obtained connected components are stitched together via carefully designed matching constraints.

A recent study of [83] introduces and Adaptive Block Coordinate optimization (ABCD) for minimizing SVD-based energies and computing inversion-free maps in \mathbb{R}^2 and \mathbb{R}^3 . ABCD is a non-linear geometric solver, based on alternating optimization of different geometric measures and on adaptive partitioning of vertices into blocks. Although this method has no theoretical guarantees to converge to an optimal solution, it can recover from a highly distorted initializations with a large fraction of inverted simplices (Fig. 1 bottom).

Finally, we present a brief discussion on the relation between injective and positively-oriented simplicial maps.

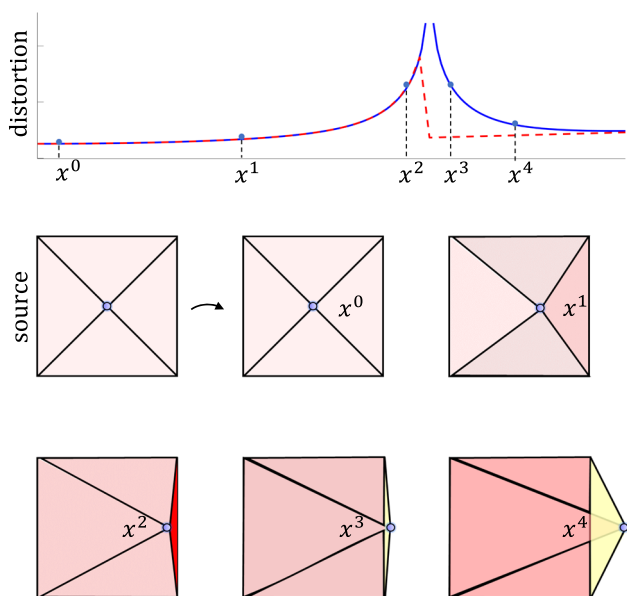


Fig. 12 Example that illustrates how inversions of triangles impact the values of barrier and non-barrier isometric distortions. At the top, we plot the two distortion energies on logarithmic scale as the source vertex is moved horizontally from the center to the right. The blue curve shows the values of the isometric barrier distortion defined by (32), and the red curve depicts the non-barrier version of (32) constructed according to [83]. The color of positively oriented triangles encodes the amount of isometric distortion. The inverted triangles are colored in yellow


5.3 Injectivity Constraints

There are two types of injective (one-to-one) simplicial maps: a globally injective and locally injective maps. Although there is close relation between one-to-one maps and orientation-preserving maps, these properties of simplicial maps are not equivalent.

A map $f \in \text{PL}(\mathcal{M}, n)$ is locally injective if it maps neighboring simplices without self-intersections. That is, f is locally one-to-one if, for any two simplices $s_1 \neq s_2$ that share common vertices, we have

$$\text{int}(\text{conv}(s'_1)) \cap \text{int}(\text{conv}(s'_2)) = \emptyset, \tag{67}$$

where s'_1, s'_2 are the corresponding target simplices under f and $\text{int}(B)$ denotes the interior of a set B . Obviously, map f is globally injective if it is locally one-to-one and $\text{conv}(s'_1) \cap \text{conv}(s'_2) = \emptyset$ for any disjoint simplices $s_1, s_2 \in \mathcal{S}$.



Satisfying (57) for a map f , alone, does not guarantee that f is locally injective. For example, if v is a vertex of a planar




Fig. 13 A positively-oriented simplicial map f that is not one-to-one. We highlight areas where the mapping is non-locally injective (b) and non-globally injective (c). The non-locally injective configuration causes high isometric distortions and artifacts in the texture mapping (a)

triangular mesh, then triangles around v can be twisted into a loop with an angle greater than 2π (see the inset). Therefore, even if the optimization begins with a globally injective map and preserves its orientation, it can produce simplicial maps that are non one-to-one, both on the local and global scales (see Fig. 13).

Nevertheless, there is a limited number of algorithms for inducing globally injective maps with low distortion on triangular and tetrahedral meshes [53,81,103]. These methods start with a globally injective map f^0 and iteratively minimize distortions in such a way that, at each iteration i , orientation of the map f^i is kept positive and there are no intersections between boundary simplices under f^i . This strategy guarantees a global bijection according to the next theorem [3,66,67]:

Theorem 5.1 *A positively oriented simplicial map $f : \text{conv}(\mathcal{M}) \rightarrow S' \subset \mathbb{R}^n$ is globally injective if the restriction of f to the boundary, $f|_{\partial \text{conv}(\mathcal{M})} : \partial \text{conv}(\mathcal{M}) \rightarrow \partial S'$, is globally bijective map.*

6 Mesh Modification Techniques

In some cases, the initial triangulation \mathcal{M} of a shape S can be modified to attain a better solution of the optimal mapping problem on S . As shown in the next two subsections, a proper modification of meshes often leads to more robust algorithms for surface parameterization and inter-surface mapping.

6.1 Parameterization and Inter-Surface Mapping

Mesh modifications allow parameterization algorithms to deal with more restrictive positional constraints. For example, algorithms of [61,99] compute bijective constrained parameterizations of surfaces by adding Steiner vertices [88]

to the original mesh and triangulating the region between a given target domain and its boundary rectangle.

In certain applications, users might only be interested in the overall shapes of the source and target domains, rather than in finding the exact locations of each mesh element. In this case, by allowing mesh refinements, linear maps (59) can be extended to more general locally injective maps of meshes onto non-convex domains. One example of such approach is the parameterization algorithm of [118] that starts with a given source mesh and a coarse triangulation of the target polygon. First, this method computes Tutte embeddings f and f' of the source mesh M and target mesh M' onto an intermediate convex domain, then it constructs the map f_\circ of M onto M' , as the composition

$$f_\circ = (f')^{-1} \circ f. \tag{68}$$

Finally, meshes M and M' are gradually refined as long as the obtained composition f_\circ remains a non-locally injective mapping.

In geometry processing, intermediate convex domains and map compositions are used in various applications, including shape matching tasks. In particular, many algorithms for bijective mapping between two surfaces (*inter-surface mapping*) start by flattening source and target meshes into planar discs and then matching these discs to find the bijection. For example, methods of [6,7,93] initialize a shape matching problem with $f^0 = f_\circ$, computed according to (68), then

these methods use different non-linear solvers to reduce isometric distortion of the inter-surface map, while positional constraints (58) are used to match between given landmark points. Likewise, mesh modifications are often used in inter-surface mapping algorithms to induce injective maps and to handle meshes with different connectivities [6,60,92].

Fig. 14 illustrates some of the above approaches to locally injective parameterization and inter-surface mapping between triangle meshes.

6.2 Global Parameterization

Many of existing techniques for locally injective mapping are limited to domains with a simple topology. Therefore, to achieve desirable results, it is often necessary to modify a given triangulation and to simplify mesh topology.

In this subsection, we present a brief introduction to the *global parameterization* – a problem of flattening meshes of a non-disc topology. This problem is closely related to linearly constrained methods for distortion minimization and locally injective mappings.

Note that most of the mapping algorithms presented so far are topology-preserving methods. Therefore, these methods can be applied only for flattening surfaces that are homeomorphic to a planar disc. However most of the real world surfaces are not homeomorphic to a disc, and therefore parameterization of such surfaces requires additional tools of global parameterization.

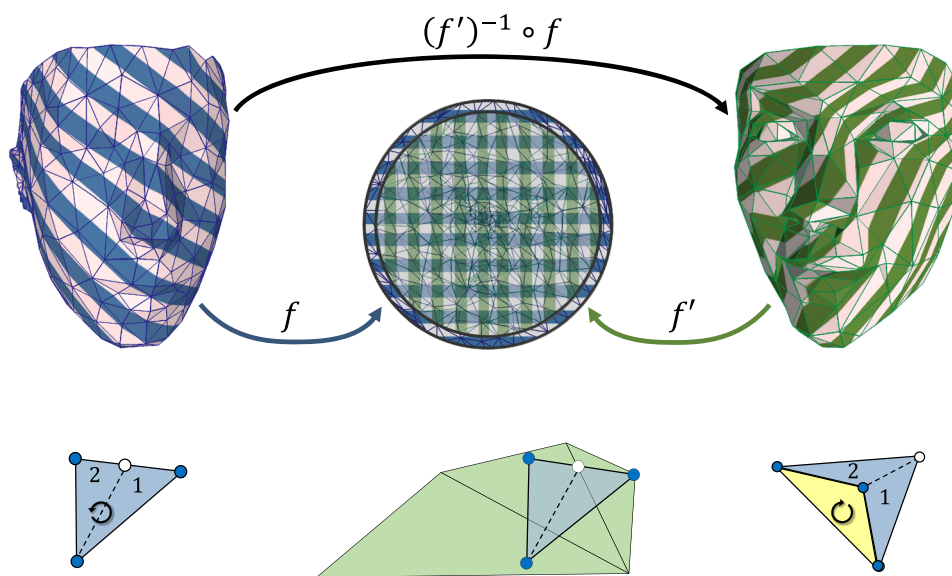


Fig. 14 Top: Injective maps f and f' of meshes $M = (S, \mathcal{V})$ and $M' = (S', \mathcal{V}')$ onto a common planar domain \mathbb{D} . Their composition $f_\circ = (f')^{-1} \circ f$ yields a locally injective mapping between the two meshes. **Bottom:** Refining the source mesh to repair inversions induced by f_\circ . We assume that M' is a planar mesh and demonstrate how the algorithm of [118] maps the selected area before and after mesh

refinements. Showing from the left to the right: a positively oriented source triangle $s \in S$; the location of the triangle $f(s)$ and its neighbors mapped by f' in \mathbb{D} ; the negatively oriented target triangle $f_\circ(s)$ (marked in yellow). If s is cut across the the dashed line, then s is split into new source triangles s_1 and s_2 which are mapped by f_\circ onto new positively oriented target triangles, shown on the right




Fig. 15 Unconstrained global parameterization (left) versus a seamless parameterization under rotation constraints (69), for $\rho_{uv} = 1$, (right)

Assume that $M = (S, \mathcal{V})$ is a triangulation of a non-disc topology surface. Typically, a global parameterization of M consists of the two primary steps: (i) cutting M into a disc-topology mesh $M_{\mathbb{D}}$; (ii) flattening of $M_{\mathbb{D}} = (S_{\mathbb{D}}, \mathcal{V}_{\mathbb{D}})$ under constrained positions of vertices that were duplicated during the first step.

There is a strong correlation between the geometry of the mesh cut and the quality of the obtained global parameterization. For example, consider parameterization of a well-cut brain surface, depicted by Fig. 1 (top), and parameterization of a similar poorly cut surface, shown in Fig. 4. Clearly, flattening of a poorly cut mesh leads to increased distortions.

In order to attain a global parameterization with low distortions, some methods for cutting meshes into a disc are based on the joint minimization of geometrical energies and quality measures of the cut [69,87,104]. Other methods for mesh cutting are based on discrete exterior calculus [77] and on constructions of guided vector fields [14,16]. For more details about mesh cutting algorithms readers are referred to: [25,44,76,77].

Clearly, the second step of parameterizing $M_{\mathbb{D}}$ is a special case of the problem of optimal mapping from \mathbb{R}^3 to \mathbb{R}^2 . In particular, the most common constraints, employed in the second stage of global parameterization, are rotation constraints aimed at aligning textures along the cut edges (*seam edges*). If $v, u \in \mathcal{V}$ are neighboring vertices of a seam edge and $v', u' \in \mathcal{V}_{\mathbb{D}} \setminus \mathcal{V}$ are their duplicates, obtained during the mesh cut, then typical rotation constraints for these vertices are expressed as the following complex equation:

$$\mathbf{x}_u - \mathbf{x}_v = e^{i\pi\rho_{uv}/2}(\mathbf{x}_{u'} - \mathbf{x}_{v'}), \quad \rho_{uv} \in \{0, 1, 2, 3\}. \quad (69)$$

If a global parameterization $f[\mathbf{x}] \in \text{PL}(M_{\mathbb{D}}, 2)$ satisfies constraints (69) for any pairs of duplicated edges (u, v) and (u', v') of $M_{\mathbb{D}}$, then f is called a *seamless parameterization* (see Fig. 15).

Note that rotation constraints are compatible with most distortion minimization techniques presented in Sect. 4.3. Therefore, seamless parameterization methods can employ

the tools from the previous sections for distortion minimization and locally injective mapping. For instance, [9,44] employ a modified Tutte embedding for initialization and the BD method [66] for ensuring a locally injective parameterization of $M_{\mathbb{D}}$. Likewise, final stages of global parameterizations [25] and [44] minimize the isometric distortion (32) via SLIM and PN solvers, respectively.

A more elaborate discussion of surface parameterization is beyond the scope of this paper. We refer readers to [49] for a survey of classical methods in global parameterization, and to [9,25,44,77,126] for the overview of more recent global parameterization methods.

7 Continuous versus Discrete Mapping

We have developed theory that works for both continuous and discrete maps, nevertheless there are some important differences between the continuous and discrete cases.

Roughly speaking, the continuous version of problem (1) is more general and admits a larger set of solutions compared to its discrete counterpart; certain types of distortion minimizers can only be realized over continuous domains, but not over discrete ones. Here we explain this inherent difference between the two settings for minimizing conformal distortions (see Definition 2.14). Global minimizers of conformal distortions are closely related to the so-called *extremal quasi-conformal maps* [115]. Recall that a map is conformal if it stretches the space equally in all directions (see Definition 2.6), i.e., for a simplicial map f , this means that for all $s \in \mathcal{S}$ singular values $\sigma_i(df_s)$ are equal. In other words, f is conformal if, for each s , linear components of df_s are uniform scaling transformations, i.e.,

$$\exists \theta : \mathcal{S} \rightarrow \mathbb{R} \text{ s.t. } \forall s \in \mathcal{S} : [df_s]_{\sim} = \left[e^{\theta(s)} I_{n \times n} \right]_{\sim}, \quad (70)$$

where \sim if the equivalence relationship from (19).

In the continuous settings, there is a fundamental difference between the plane and higher dimensions. According to Riemman mapping theorem, any pair of simply-connected domains in the plane can be mapped conformally. In contrast, Liouville’s theorem [59] states that conformal maps in dimensions $n \geq 3$ are restricted to the limited set of Möbius transformations, consisting of the compositions of rigid transformations, uniform scaling and inversions on a sphere.

The notion of continuous conformal maps is also well studied in Riemannian settings, where the conformality is formulated for both Riemannian metrics and for mapping between Riemannian manifolds according to the following definitions:

Definition 7.1 (Riemannian conformal metric) Two Riemannian metrics g_1 and g_2 of a manifold \mathcal{M} are called




Fig. 16 From left to right a properly connected simplicial complex, a not properly connected simplicial complex and its image under a conformal simplicial map. Note that Theorem 7.1 does not apply to a path-connected mesh \mathcal{M} if \mathcal{M} is not properly connected

conformally equivalent if

$$\exists \theta : \mathcal{M} \rightarrow \mathbb{R} \text{ s. t. } g_1 = e^\theta g_2. \tag{71}$$

Definition 7.2 (Riemannian conformal map) A diffeomorphism $f : (\mathcal{M}, g) \rightarrow (\mathcal{N}, h)$ between two Riemannian manifolds is called *conformal* if the pulled back metric f^*h , induced by f on \mathcal{M} , is conformally equivalent to g .

The function e^θ in (71) and (70) is called a *conformal factor*. In both the discrete and continuous settings, conformal factors can be interpreted as uniform scaling of source domain caused by a map.

According to the above definitions, each conformal map f can be associated with its conformal factor, defined over the source domain. This view leads to the interpretation of the conformal mapping as a transformation that preserves shape of spheres, which gives an alternative definition of a conformal map.

In the continuous scenario, spheres can be infinitesimally small, leading to a high degree of freedom in setting conformal factors. However, in the case of triangulated domains, conformal factors are set per simplex and so that the conformal factors coincide on common faces of neighboring simplices. Thus, once $\theta(s_0)$ is set on a simplex s_0 , it automatically predefines θ on all simplices s_k that are face-to-face connected to s_0 . By the face-to-face connection, we mean that there is a path of simplices s_0, s_1, \dots, s_k such that each pair $\{s_{j-1}, s_j\}$ shares a common face. We call $\mathcal{M} = (\mathcal{S}, \mathcal{V})$ a *properly connected* complex if all of its simplices are face-to-face connected (see Fig. 16 for the illustration). A conformal map f of a properly connected mesh can be specified, up to a rigid motions of target simplices, by defining a single linear component df_{s_0} . This restriction of simplicial maps is formalized below.

Theorem 7.1 Let $f : \text{conv}(\mathcal{M}) \rightarrow \mathbb{R}^d$ be a simplicial mapping of n -dimensional properly connected simplicial complex $\mathcal{M} = (\mathcal{S}, \mathcal{V}, \mathbf{y})$, embedded in \mathbb{R}^m , $m \geq n$ and $d \geq n$. If f is conformal, then, for each $s \in \mathcal{S}$, f_s is a composition of a rigid transformation from \mathbb{R}^m to \mathbb{R}^d and the same uniform scaling of \mathbb{R}^d .

Proof Pick the conformal factor $e_0 = e^{\theta(s_0)}$ on some simplex s_0 . Then, $e^{\theta(s)} = e_0$ for all the neighbors s of s_0 , since otherwise \mathcal{M} is not properly embedded. Furthermore, the latter equality holds for any $s \in \mathcal{S}$, since \mathcal{M} is connected. Consequently, (46) and (70) imply

$$\forall s \in \mathcal{S} : \tilde{f}_s = \psi_s \circ f_s \circ \phi_s^{-1}, \tag{72}$$

$$[d\tilde{f}_s]_{\sim} = [\text{diag}(e_0, \dots, e_0)]_{\sim}, \tag{73}$$

where ψ_s and ϕ_s are transition maps, set per simplex according to (46). Assume w.l.o.g. that ϕ_s and \tilde{f}_s are linear transformations, then

$$f_s = \psi_s^{-1} \circ \text{diag}(e_0, \dots, e_0) \circ \phi_s \tag{74}$$

$$= \psi_s^{-1} \circ \phi_s \circ \text{diag}(e_0, \dots, e_0), \tag{75}$$

where $(\psi_s^{-1} \circ \phi_s)$, $s \in \mathcal{S}$, are rigid transformations and $\text{diag}(e_0, \dots, e_0)$ is the uniform scaling, as stated by the theorem. \square

Consequently, in contrast with the continuous scenario, discrete conformal maps constitute a very restricted family of maps, regardless of whether the maps are defined in the plane or in higher dimension.

From the conformal geometry viewpoint, patterns of spheres are more convenient structures than simplicial complexes because spheres are the basic invariants of conformal maps. In particular, if a surface is divided into a pattern of circles, then a continuous conformal map on that surface can be discretized by specifying a Möbius transformation per each circle. This approach, known as a *circle parking* metric, is used to allow more flexibility in the underlying conformal structure of discrete mappings.

In the plane, Möbius transformation are conformal maps that transform circles into circles.¹² These transformations preserve the so-called four point *cross-ratio*

$$\text{CR}(\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3, \mathbf{r}_4) \triangleq \frac{(\mathbf{r}_3 - \mathbf{r}_1)(\mathbf{r}_4 - \mathbf{r}_2)}{(\mathbf{r}_3 - \mathbf{r}_2)(\mathbf{r}_4 - \mathbf{r}_1)}, \tag{76}$$

defined for complex numbers $\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3, \mathbf{r}_4 \in \mathbb{C}$. In particular, (76) implies that a unique Möbius transformations $f(\mathbf{r}) : \mathbb{C} \rightarrow \mathbb{C}$ can be defined using the cross-ratio equation,

$$\text{CR}(\mathbf{r}, \mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3) = \text{CR}(f(\mathbf{r}), f(\mathbf{r}_1), f(\mathbf{r}_2), f(\mathbf{r}_3)), \tag{77}$$

that specifies the given correspondence between source points \mathbf{r}_j and target points $f(\mathbf{r}_j)$, for $j = 1, 2, 3$. This property is used in various applications for planar deformations and shape analysis. For example, Möbius transformations

¹² Here we consider the generalized notation according to which lines are circles with infinite radius.

between critical points of the average geodesic distance are used in [58] for a global symmetry analysis of surfaces. Lipman et al. [68] extend (77) to an analytic formula for computing extremal quasi-conformal mappings between quadruplets.

The circle parking approach and the related preservation of the cross-ratio leads to the “piecewise Möbius paradigm” for discrete conformal mapping. According to this paradigm, conformal maps are approximated by As-Möbius-As-Possible (AMAP) transformations of adjacent triangles. This leads to two alternative definitions of discrete conformality. According to the first definition a map f of a triangle mesh is conformal if it preserves intersecting angles of triangle circumcircles [62]; whereas according to the second definition f is conformal if it preserves the absolute values of cross-ratios (76), computed per each pair of adjacent triangles [105]. Notably, the cross-ratio definition of the conformality can be extended, with certain limitations, to maps $f : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ by representing 3D coordinates via imaginary quaternions $\text{Img}(\mathbb{H})$ (see [22]) and by considering a quaternionic cross-ratio [111,112], i.e., the quantity $\text{CR}(\mathbf{r}_1, \dots, \mathbf{r}_4)$ defined according to (76) for points $\mathbf{r}_1, \dots, \mathbf{r}_4 \in \text{Img}(\mathbb{H})$.

Conformal maps of circle patterns can be encoded by specifying edge-based conformal factors on triangular meshes. In this case, map f is a conformal map if it scales uniformly half-edges sharing the same vertex (see illustration in Fig. 17). This metric approach is employed in computations of conformal maps based on Ricci flow and solutions of the Beltrami equation [19,122].

Although circle parking metric naturally extends to non-Euclidean geometries, conformal factors of this metric are defined implicitly via the length of simplicial complex edges or radii of circumscribed circles. Therefore, unlike the finite element formulation, in the circle parking metric, conformal maps are constructed in two separate steps: first, conformal factors are computed, then the obtained metric is embedded into the target space. The embedding step can be formulated as optimization problem (1) with respect to isometric distortion; it brings our discussion back to indirect methods and other distortion optimization approaches, covered in Sect. 4.3.

While simplicial conformal maps are prescribed by setting a single conformal factor, a general map $f \in \text{PL}(\mathbf{M}, n)$ is characterized by its singular values $\sigma_1(df_s), \dots, \sigma_n(df_s)$, computed over a subset of simplices in \mathcal{S} . However, if simplicial map f possesses a low conformal distortion, then, in practical terms, f can be described by a set of *approximate conformal factors*. In particular, if $f \in \text{PL}(\mathbf{M}, n)$ is a quasi-conformal map, approximating a continuous conformal map f_{cont} , then f induces a low conformal distortion, so that its components $f_s, s \in \mathcal{S}$, are close to be similarity transformations. In this case, conformal factors of f_{cont} can




Fig. 17 A circle packing metric illustrated for an approximate conformal parameterization, computed by the BFF method [94]. According to Theorem 7.1, the obtained parameterization f is a quasi-conformal mapping. Therefore, some planar circles are mapped by f^{-1} into ellipses

be approximated by computing average singular values of f on simplices $s \in \mathcal{S}$:

$$\mathcal{D}_{\text{CF}}(df_s) \triangleq (\sigma_1(df_s) + \dots + \sigma_n(df_s))/n. \tag{78}$$

As shown by Fig. 1, a histogram of distortions (78) can be employed as a map descriptor or as a shape signature for a collection of 3D objects, mapped onto a common domain [8].

8 Convex Analysis of Distortions

The main challenge in optimal mapping of a mesh $\mathbf{M} = (\mathcal{S}, \mathcal{V})$ stems from the fact that, in most cases, objective measure $E(\mathbf{x})$ and a set of feasible target coordinates,

$$X_f(\mathbf{M}, n) \triangleq \{\mathbf{x} \in \mathbb{R}^{n|\mathcal{V}|} \mid \forall s \in \mathcal{S} : \det df_s[\mathbf{x}] > 0\}, \tag{79}$$

are both non-convex. Therefore, to obtain a convex approximation of the problem one need to modify both the energy E and the set X_f . Usually, finding a convex subset of X_f requires to remove a significant part of the feasible set, as the shape of X_f is very irregular. Despite the highly-non convex structure of (79), certain geometric measures can be realized by simple distortion energies that are convex in $\mathbf{x} \in \mathbb{R}^{n|\mathcal{V}|}$, over the entire set of target coordinates. If $E(\mathbf{x})$ is a convex distortion energy and $\text{conv}^i(X_f) \subset X_f$ is a convex subset of (79) constructed at iteration i , then optimizing $E(\mathbf{x})$ in $\text{conv}^i(X_f)$ avoids the need for the energy approximation step (64). Therefore, minimizing a convex distortion leads to better convergence than minimizing a similar non-convex energy. In particular, in a convex scenario, the exact Hessian of $E(\mathbf{x})$ can be used in (63) and (64) for Newton optimization. Furthermore, in certain problems, such as harmonic mapping into a planar disc or minimization of TLC,




Fig. 18 Examples of convex distortion energies, introduced in Sect. 8. From the left to right are depicted: source tetrahedral mesh with fixed position of endpoint vertices (anchors) and resulting minimization of Dirichlet energy (25) and of symmetric gauge distortions induced by L_2

and L_3 norms, respectively. All optimization problems are initialized by the identity map and the solutions are obtained via BCQN solver [121]

orientation requirements (57) can be substituted by setting proper boundary constraints. In such cases, computing a local minimizer of (56) under constraints (58) is guaranteed to achieve the global minimum of a convex measure $E(\mathbf{x})$.

For example, we can penalize ‘stretching’ of the Euclidean space by convex distortion measures, including the Dirichlet energy (25) and some of its variants [72]. To the best of our knowledge, there is no larger family of SVD-based distortions that are proven to be convex in \mathbf{x} and that are also used in practical applications.

In view of the above considerations, the first goal of this section is to introduce a new family of convex distortions. Our second goal is to introduce a convex analysis of distortions and to identify necessary and sufficient conditions for distortion measures to be convex functions. This gives us criteria by which we can formally prove that the vast majority of existing distortion energies are non-convex in \mathbf{x} .

We begin by expanding analysis of the relevant mapping spaces. Due to the determinant constraints (57), the space $X_f(\mathbf{M}, n)$ constitutes an open and non-convex subset of $\mathbb{R}^{n|\mathcal{V}|}$. Moreover, even the smaller set of non-degenerate simplicial maps

$$\text{Diff}(\mathbf{M}, n) \triangleq \{f \in \text{PL}(\mathbf{M}, n) \mid \det df_s \neq 0, s \in \mathcal{S}\}$$

and its continuous counterpart $\text{Diff}(\mathbb{R}^n)$ are both highly non-convex. Here by ‘highly non-convex’ we mean that, at any point, any large-enough subspace containing this point is non-convex. Indeed, maps in the resulting space can be identified with their Jacobians. Thus, the set of non-singular matrices $J \in \text{GL}(\mathbb{R}, n)$, characterizes the family of non-degenerate maps by prescribing the Jacobian at each point, or over each simplex. The set $\text{GL}(\mathbb{R}, n)$ is not closed under addition, and the radius of the maximal convex subset around J of $\text{GL}(\mathbb{R}, n)$, is $\sigma_n(J)$, i.e., the smallest singular value of J . Consequently, for a given $f \in \text{Diff}(\mathbb{R}^n)$, the diameter of a

maximal convex-subset of f is a function of the chosen norm on $\text{Diff}(\mathbb{R}^n)$ and the smallest singular value $\sigma_n(df_r)$, $r \in \mathcal{S}$ (in the discrete case, $\sigma_n = \sigma_n(df_s)$, $s \in \mathcal{S}$).

Due to the non-convex structure of the mapping spaces, most of the existing methods for solving (56) are based on iterative optimization algorithms — each iteration of these algorithms modifies target coordinates in a small convex neighborhood of $\mathbf{x}^i \in \text{conv}^i(X_f(\mathbf{M}, n))$, where results are guaranteed to satisfy (54). As pointed out in Sect. 4.3, there are many methods for building subsets $\text{conv}^i(X_f)$, including representations of BD spaces by cones and polytopes, computations of line search intervals for non-linear solvers, and more. Since it is impractical to analyze distortion convexity for every possible choice of the subset $\text{conv}^i(X_f)$, we do not restrict our analysis to a specific convexification of the mapping space. Thus, we treat distortion energies as functions of $\mathbf{x} \in \mathbb{R}^{n|\mathcal{V}|}$.

One the one hand, this approach allow a more straightforward analysis of distortions. On the other hand, having no restrictions on target coordinates \mathbf{x} can produce points where the linear map $df_s[\mathbf{x}]$ has zero singular values. To overcome this limitation, we assume that all the distortions measures, considered in this section, are normalized and extensible to $\overline{\mathbb{L}^n}$, the closure of \mathbb{L}^n in \mathbb{R}^n . That is, we consider normalized distortions $\mathcal{D}(\sigma)$ that can be extended to the function $\overline{\mathcal{D}} : \overline{\mathbb{L}^n} \rightarrow [0, \infty]$, as follow:

$$\overline{\mathcal{D}}(\sigma) \triangleq \begin{cases} \mathcal{D}(\sigma) & \sigma \in \mathbb{L}^n \\ \lim_{\sigma' \rightarrow \sigma} \mathcal{D}(\sigma') & \sigma \in \partial \overline{\mathbb{L}^n}. \end{cases} \quad (80)$$

If the limit of \mathcal{D} exists in (80) for every $\sigma \in \partial \overline{\mathbb{L}^n}$, then $\overline{\mathcal{D}}$ is called an *extended distortion* of \mathcal{D} . It is easy to show that extended distortions satisfy the essential properties of Definition 2.11, for differentiable maps $f \in C^1(\mathbb{R}^n)$. If $\overline{\mathcal{D}}$ is the extended distortion of \mathcal{D} , then $\overline{E} = E_{\overline{\mathcal{D}}}(\mathbf{x})$, defined according to (47), is the extension of the energy $E_{\mathcal{D}}(\mathbf{x})$ from

$\{\mathbf{x} \in \mathbb{R}^{n|\mathcal{V}|} \mid f[\mathbf{x}] \in \text{Diff}(M, n)\}$ to the arbitrary target coordinates $\mathbf{x} \in \mathbb{R}^{n|\mathcal{V}|}$.

Since \overline{E} and \overline{D} can be infinite, we use the *epigraph definition* of the function convexity. That is, $\overline{E}(\mathbf{x})$ is a convex function of \mathbf{x} if its epigraph, denoted by $\text{epi } \overline{E}$, is a convex subset of $\mathbb{R}^{n|\mathcal{V}|}$ (see the example at the bottom of Fig. 19). We use the same notion of the convexity for \overline{D} with respect to arguments $J \in \mathbb{R}^{n \times n}$ and $\sigma \in \overline{\mathbb{L}}^n$. Likewise, we say that $\mathcal{D}(J)$ and $\mathcal{D}(\sigma)$ are convex if \overline{D} is well-defined convex function of $J \in \mathbb{R}^{n \times n}$ and of $\sigma \in \overline{\mathbb{L}}^n$, respectively.

As shown by (52), there is a linear transformation that identifies Jacobians df_s with the target coordinates \mathbf{x} . We can, thus, detect which distortion measures are convex with respect to \mathbf{x} by analyzing the convexity of distortions as if they are functions defined over $n \times n$ Jacobian matrices. Hence we extend existing convex measures by introducing a new family of distortions and proving that elements of this family are convex functions in the Jacobian.

Due to the triangle inequality, any matrix norm $\|J\|$, $J \in \mathbb{R}^{n \times n}$ is a convex function of $J \in \mathbb{R}^{n \times n}$. Consequently, distortion \overline{D} is a convex measure in $J \in \mathbb{R}^{n \times n}$, if there is a matrix norm $\|\cdot\|$ such that $\overline{D}(J) = \|J\|$ for each J . These observations are formalized by the following lemma:

Lemma 8.1 *Let $\|\cdot\|$ be a matrix norm in $\mathbb{R}^{n \times n}$ and define*

$$\mathcal{D} : \text{GL}(\mathbb{R}, n) \rightarrow \mathbb{R}, \quad \mathcal{D}(J) \triangleq \|J\|. \tag{81}$$

If $\|\cdot\|$ is a unitary invariant norm, that is,

$$\|RJ\| = \|JR\| = \|J\|, \quad \forall J, R \in \mathbb{R}^{n \times n}, |\det R| = 1, \tag{82}$$

then, \mathcal{D} is a first-order distortion and $\overline{D}(J)$ is convex in $J \in \mathbb{R}^{n \times n}$.

Proof The proof is immediate: due to the norm convexity and norm continuity, $\overline{D}(J) = \|J\|$ is a convex function of matrices; \mathcal{D} satisfies first order precision (Definition 2.11) by its definition, and (9) is met by our assumption on the unitary invariance of $\|\cdot\|$. \square

Lemma 8.1 implies that each unitary invariant matrix norm $\|\cdot\|$ defines, via (81), a convex distortion density $\mathcal{D}(f, \mathbf{r}) = \|df_{\mathbf{r}}\|$. Indeed, if $\|\cdot\|$ is unitary invariant, then SVD of $J \in \mathbb{R}^{n \times n}$ implies that $\|J\|$ can be expressed as a function of J 's singular values

$$\|J\| = g(\sigma_1(A), \dots, \sigma_n(A)), \tag{83}$$

and, therefore, the restriction of g to $\text{GL}(\mathbb{R}, n)$ is a distortion measure, according to Theorem 2.1.

According to matrix analysis, a function g that satisfies (81) with a unitary invariant norm $\|\cdot\|$ is called a *symmetric gauge function*. There is a number of equivalent ways to

define symmetric gauge functions. We adopt the following definition [47]:

Definition 8.1 A function $g : \mathbb{R}^n \rightarrow (-\infty, \infty]$ is called *absolutely permutation-symmetric* if for any $\mathbf{r} \in \mathbb{R}^n$ and permutation P of \mathbb{R}^n

$$g(|\mathbf{r}|) = g(\mathbf{r}) = g(P\mathbf{r}).$$

An absolutely permutation-symmetric function g is called a *symmetric gauge function* if there exist a vector norm $\|\cdot\|$ such that, for any $\mathbf{r} \in \mathbb{R}^n$, $g(\mathbf{r}) = \|\mathbf{r}\|$.

In particular, all unitary invariant norms are characterized by symmetric gauge function via the following theorem [47, pp. 438-439]:

Theorem 8.1 *A matrix norm $\|J\|$, $J \in \mathbb{R}^{n \times n}$ is unitary invariant norm iff there is a symmetric gauge function g , such that $\|J\| = g(\sigma_1(J), \dots, \sigma_n(J))$.*

Consequently, according to Theorem 8.1 and Lemma 8.1, the following distortion measures are convex functions of Jacobian matrices:

$$\{\mathcal{D}_g : J \mapsto g(\sigma_1(J), \dots, \sigma_n(J)) \mid g - \text{symmetric gauge}\}. \tag{84}$$

We refer to the set (84) as to the *symmetric gauge distortions*. This subset includes the following well known unitary invariant norms [48, pp. 465-466]:

- L_p -norms:

$$\mathcal{D}_{L_p}(\sigma) \triangleq \|\sigma\|_p = (\sigma_1^p + \dots + \sigma_n^p)^{1/p}, \quad 1 \leq p < \infty.$$
- The spectral norm:

$$\mathcal{D}_{\text{spec}}(df_{\mathbf{r}}) \triangleq \|df_{\mathbf{r}}\|_2 = \sigma_1(df_{\mathbf{r}}).$$
- Ky Fan k -norms:

$$\mathcal{D}_{[k]}(\sigma) \triangleq \sigma_1 + \dots + \sigma_k, \quad k = 1, 2, \dots, n.$$

Obviously, we can extend (84) to a larger family of convex distortions, by applying basic operations that preserve the convexity, such as raising \mathcal{D}_g into a positive power, or using a convex combination (38) of symmetric gauge distortions.

As we have previously mentioned, if distortion \mathcal{D} is convex with respect to Jacobian matrix, then, in the discrete case, $\mathcal{D}(f_s[\mathbf{x}])$, $s \in \mathcal{S}$, is convex as a function of target vertex coordinates. Therefore, energies (47) of symmetric gauge distortions (84) and of their variants are convex in \mathbf{x} .




Fig. 19 Illustration of the convex analysis of Theorem 8.2 in 2D. **Right:** Target triangles induced by the interpolated coordinates $\mathbf{x}(t)$, computed according to (85). Triangles are colored in the same way as in Fig. 12. **Left:** We plot energies $E_{\mathcal{D}_{ARAP}}(\mathbf{x}(t))$ and $\log(E_{\mathcal{D}_{SD}}(\mathbf{x}(t)))$ as functions of t , where \mathcal{D}_{ARAP} and \mathcal{D}_{SD} are the distortion measures (33) and (32). At the bottom-left, we highlight the epigraph of symmetric Dirichlet energy, shown on a logarithmic scale

In Fig. 18, we illustrate volumetric mappings that minimize symmetric gauge distortions.

In a certain way, symmetric gauge distortions are generalizations of the Dirichlet energy (25) that assesses the degree to which \mathbb{R}^n is stretched under f . However, these distortions could not be used for estimating other geometric measures such as length constructions, angle deviations and more. Notably, as shown by the next theorem, convex distortions cannot be used to assess some of the fundamental geometric measures:

Theorem 8.2 *Distortion energy $E_{\overline{\mathcal{D}}}(\mathbf{x})$ is neither convex nor concave in \mathbf{x} if $\overline{\mathcal{D}}$ is well-defined extension of distortion \mathcal{D} and any of the following statements is true:*

1. \mathcal{D} is an isometric distortion;
2. \mathcal{D} is a conformal distortion;
3. \mathcal{D} is a (unsigned) volume distortion;
4. \mathcal{D} is a normalized barrier distortion, i.e., $\mathcal{D}(\sigma)$ holds the normalization, the bottom and top barrier properties of Corollary 2.1.

By the isometric, conformal or volume distortion we mean that $J \in GL(\mathbb{R}, n)$ is a global minimum of $\mathcal{D}(J)$ iff the linear map J is isometric, conformal or (unsigned) volume-preserving map, respectively.

Proof We provide a constructive proof for an arbitrary isometric distortion \mathcal{D} . Fig. 19 illustrates our proof for two-dimensional distortions, defined by (33) and (32). For

non-isometric measures, repeat the same proof with a conformal or volume distortion \mathcal{D} .

Consider the complex of a single simplex $M = (\{s\}, \mathcal{V}, \mathbf{y})$, embedded in \mathbb{R}^n , and the following target coordinates:

$$\mathbf{x}^{(t)} = (1 - t)\mathbf{y} + t\overline{\mathbf{y}}, \quad t \in \mathbb{R}, \tag{85}$$

where $\overline{\mathbf{y}}$ denotes the reflection of \mathbf{y} in the plane $\{0\} \times \mathbb{R}^{n-1}$. Observe the value of $\overline{E}(f^{(t)}) = \text{Vol}(s)\overline{\mathcal{D}}(f^{(t)})$ induced by simplicial map $f^{(t)} = f[\mathbf{x}^{(t)}]$. According to our assumption, the identity map $f^{(0)}$ and the reflection $f^{(1)}$ equal to the global minimum \overline{E}^* of an isometric distortion energy \overline{E} , and $\overline{E}(f^{(t)}) = \overline{E}(\mathbf{x}^{(t)}) > \overline{E}^*$ for any $t \notin \{0, 1\}$ because for any such t the map $f^{(t)}$ is non-isometric (also non-conformal and non-volume-preserving). Denote by $\overline{E}^{[t_0, t_1]}$ the following segment of a line:

$$\overline{E}^{[t_0, t_1]} = \{(1 - t)\overline{E}(\mathbf{x}^{(t_0)}) + t\overline{E}(\mathbf{x}^{(t_1)}) | t \in [0, 1]\}. \tag{86}$$

Then, the segment $\overline{E}^{[0, 1]}$ is located below the epigraph $\text{epi } \overline{E}(f[\mathbf{x}])$, while the segment $\overline{E}^{[-1, 2]}$ resides above the points $\overline{E}(\mathbf{x}^{(t)})$, $t = 0, 1$. This immediately implies that \overline{E} is non-convex and non-concave in $\mathbf{x} \in \mathbb{R}^n \setminus \mathcal{V}$.

Similar proof applies to a normalized barrier distortion: thus we conclude that for

$$\tau = \inf \{t \in (0.5, 1] | \overline{E}(\mathbf{x}^{(t)}) = \overline{E}^*\},$$

$\overline{E}^{[1-\tau, \tau]}$ is located below $\text{epi } \overline{E}$ and $\tau > 0.5$, whereas, for a large enough $T \in (1, \infty)$, the segment $\overline{E}^{[1-T, T]}$ is placed above points $\overline{E}(\mathbf{x}^{(1-\tau)})$ and $\overline{E}(\mathbf{x}^{(\tau)})$. \square

Our next task is to provide a simple, yet general, criterion for the convexity of extended distortions. This criterion is based on the following property of unitary invariant matrix functions [11,64]:

Proposition 8.1 *Suppose that $g : \mathbb{R}^n \rightarrow (-\infty, \infty]$ is continuous and absolutely permutation-symmetric function. Then, the function $g(\sigma(J))$ of matrices $J \in \mathbb{R}^{n \times n}$ is convex iff g is convex.*

This proposition implies the following theorem:

Theorem 8.3 *Assume that $\overline{\mathcal{D}}$ is the extended distortion of a continuous distortion measure \mathcal{D} and denote by \mathbf{q}^\downarrow the vector of components of $\mathbf{q} \in \mathbb{R}^n$, sorted in descending order. Then, $\overline{\mathcal{D}}(J)$ is convex in $J \in \mathbb{R}^{n \times n}$ iff $\overline{\mathcal{D}}(|\mathbf{q}^\downarrow|)$ is convex in $\mathbf{q} \in \mathbb{R}^n$. Furthermore, if $\overline{\mathcal{D}}(|\mathbf{q}^\downarrow|)$ is convex in \mathbf{q} , then the distortion energy $E_{\overline{\mathcal{D}}}$ is convex in \mathbf{x} .*

Proof Define the following vector function:

$$g(\mathbf{q}) = \overline{\mathcal{D}}(|\mathbf{q}^\downarrow|).$$

Clearly, g is continuous and absolutely permutation-symmetric function and $\overline{\mathcal{D}}(J) = g(\sigma(J))$. Therefore, the first statement of the theorem is proved by applying Proposition 8.1 for the function g . The second statement is true because, for any simplicial complex $(\mathcal{S}, \mathcal{V})$ and any simplex $s \in \mathcal{S}$, $\mathbf{x} \mapsto df_s[\mathbf{x}]$ is a linear map from $\mathbb{R}^{n|\mathcal{V}|}$ to $\mathbb{R}^{n \times n}$. \square

According to Theorems 8.2 and 8.3, symmetric gauge distortions and their variants are the only convex distortion measures introduced so far in our paper. Noteworthy, some distortions $\mathcal{D}(\sigma)$ could not be extended to the set of non-negative singular values and therefore our convex analysis is not applicable to these distortions. In particular, distortions (26), (27) and (29) contain singular value ratios and thereby are not extensible to $\overline{\mathbb{L}^n}$. Although these distortions are non-convex in a general case, some of them are proven to be convex when restricted to small subsets of X_f . For example, if $f[\mathbf{x}]$ is inversion-free flattening of a triangle mesh and \mathbf{x}_v is a target coordinates of a single vertex $v \in \mathcal{V}$, then $E_{\text{MIPS}_{2D}}(\mathbf{x}_v)$ is convex in the interior of the one-ring of $f(v)$ [45].

Finally, to complete our discussion on the convexity, we would like to notice that certain non-convex energies $E(\mathbf{x})$ become convex when considered as functions of other mesh parameters. For instance, the symmetric Dirichlet energy (32) is convex with respect to edge length squares (ELS). This property of ELS-based energies is employed in a number of recent studies on shape interpolations [1] and surface parameterization in a metric domain [21].

9 Multi-Resolution Invariance of Distortions

Our approach to optimizing geometric distortions is built upon a piecewise linear approximation of the real-world continuous deformations. However, there exist infinitely many tessellations of the same proper domain S , leading to an infinite number of possibilities for representing continuous deformations of S . Thus, one should formally prove that problem (53) is well defined in the sense that “equivalent” simplicial complices induce equivalent simplicial maps and equivalent minimizers of these maps. First, to formulate the concept of equivalent simplicial complices, we define the following notion of simplex refinement:


Definition 9.1 Let \mathcal{S} , \mathcal{C} and \mathcal{V} , \mathcal{U} be two simplex sets and two vertex sets, respectively, and let $\mathcal{S} = (\mathcal{S}, \mathcal{V}, \Psi_{\mathcal{S}})$ and $\mathcal{C} = (\mathcal{C}, \mathcal{U}, \Psi_{\mathcal{C}})$ be two n -dimensional simplicial complices of these sets, such that their interiors are locally embedded into \mathbb{R}^m , $m \geq n$, by functions $\Phi_{\mathcal{S}} : \mathcal{V} \rightarrow \mathbb{R}^m$ and $\Psi_{\mathcal{C}} : \mathcal{U} \rightarrow \mathbb{R}^m$. For $s \in \mathcal{S}$, denote by $\Psi_{\mathcal{S}}(s)$ and $\text{Vol}(\Psi_{\mathcal{S}}(s))$ the convex hull of vertices $\{\Psi_{\mathcal{S}}(v) | v \in s\}$ and the n -dimensional volume of the convex hull, respectively. By $\Psi_{\mathcal{C}}(c)$ and $\text{Vol}(\Psi_{\mathcal{C}}(c))$

we refer to the same notations, defined for a simplex $c \in \mathcal{C}$ and for the function $\Psi_{\mathcal{C}}$. We call \mathcal{C} a refinement of \mathcal{S} if

$$\sum_{c \in \mathcal{C}} |\text{Vol}(\Psi_{\mathcal{C}}(c))| = \sum_{s \in \mathcal{S}} |\text{Vol}(\Psi_{\mathcal{S}}(s))| , \tag{87}$$

and for each $s \in \mathcal{S}$ there exists a finite set of simplices $c_1, \dots, c_k \in \mathcal{C}$, such that

$$\Psi_{\mathcal{C}}(c_1) \cup \Psi_{\mathcal{C}}(c_2) \cdots \cup \Psi_{\mathcal{C}}(c_k) = \Psi_{\mathcal{S}}(s) . \tag{88}$$



Next, we extend the notion of the refinement to simplicial maps as follows:

Definition 9.2 Let g and f be simplicial maps defined over simplicial complices $\mathcal{C} = (\mathcal{C}, \mathcal{U}, \Psi_{\mathcal{C}})$ and $\mathcal{S} = (\mathcal{S}, \mathcal{V}, \Psi_{\mathcal{S}})$, respectively (see the inset). We call g a refinement of f if \mathcal{C} is a refinement of \mathcal{S} and if

$$g_c(\Psi_{\mathcal{C}}(u)) = f_s(\Psi_{\mathcal{C}}(u)) , \tag{89}$$

for any $u \in c \in \mathcal{C}$ and any $s \in \mathcal{S}$ containing vertices of c .

We use Definition 9.2 to formulate the multi-resolution property of distortion measures in the following theorem:

Theorem 9.1 Let $\mathcal{D}(\sigma) : \mathbb{L}^n \rightarrow \mathbb{R}$ be a distortion measure and $\mathcal{C} = (\mathcal{C}, \mathcal{U}, \Psi_{\mathcal{C}})$ be a refinement of $\mathcal{S} = (\mathcal{S}, \mathcal{V}, \Psi_{\mathcal{S}})$. Assume that simplicial map f of \mathcal{C} is a refinement of simplicial map g of \mathcal{S} , then

$$E_{\mathcal{D}}(f) = E_{\mathcal{D}}(g) ,$$

where $E_{\mathcal{D}}$ is the distortion energy (47) of f and g , computed with respect to \mathcal{D} and volume weights $w(s) = \text{Vol}(\text{conv}(s))$.

Theorem 9.1 follows from the two facts: (i) piecewise affine approximation of a piecewise affine map f is the same map f ; (ii) refinements preserve the total simplex volume. The formal proof of Theorem 9.1 is presented below:

Proof of Theorem 9.1 Each simplex $s \in \mathcal{S}$ can be decomposed into simplices $c_1, \dots, c_k \in \mathcal{C}$, such that these simplices satisfy (88) and

$$\sum_{i=1}^k \text{Vol}(\Psi_{\mathcal{C}}(c_i)) = \text{Vol}(\Psi_{\mathcal{S}}(s)) . \tag{90}$$

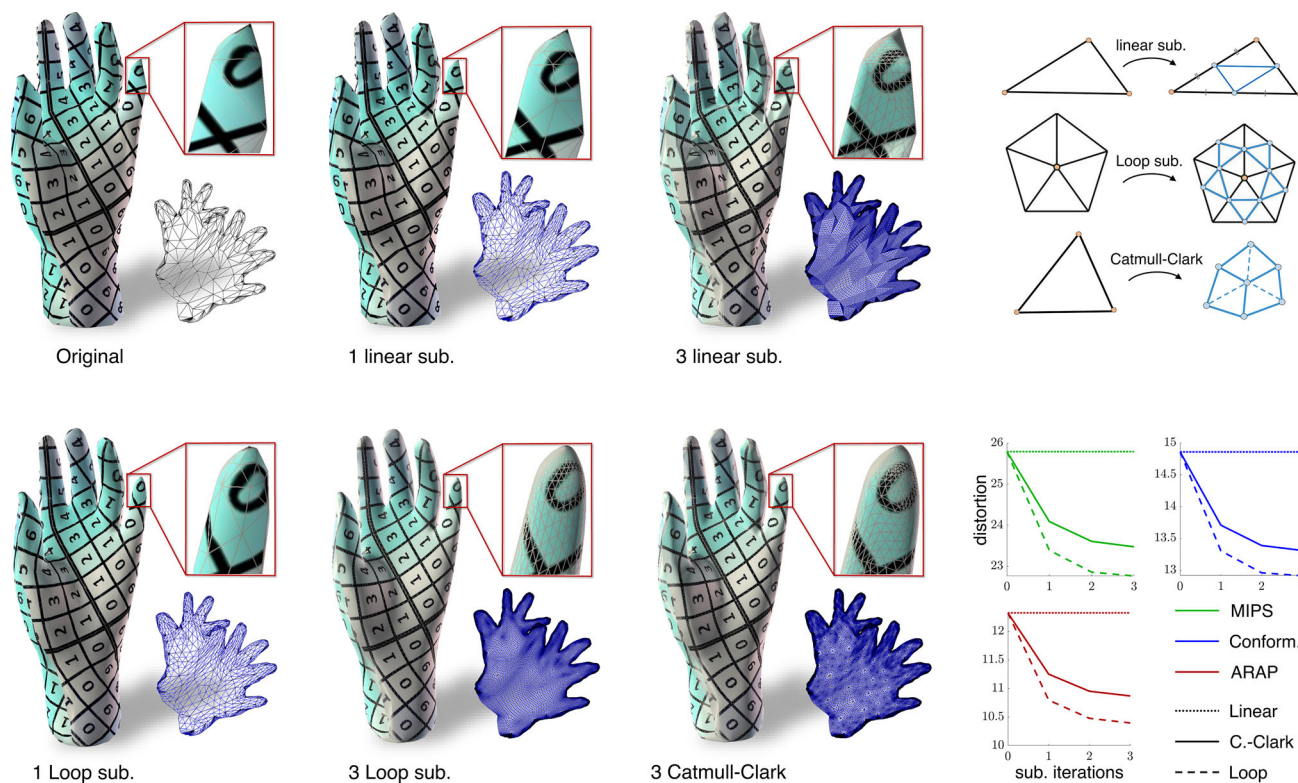


Fig. 20 Using subdivision schemes to refine the initial parameterization of the triangular mesh (top left). We use the following schemes: linear subdivision, which is an example of simplicial map refinement; Loop subdivision scheme; triangulated Catmull-Clark scheme, i.e., Catmull-Clark scheme in which polygons are triangulated at each iteration. We

plot distortions (26), (29) and (33) as functions of a number of subdivision iterations (bottom right). Note that we subdivide target and source domains by the same scheme to keep the right correspondence between their simplices

Equality (90) follows from (87) and from our underlying assumption that embedded simplices are consistently oriented. Since g is a refinement of f , for each $c = c_k$, these mappings satisfy $n + 1$ equations (89) with vertices $u_1, \dots, u_{n+1} \in c$. The resultant system of $n + 1$ equations defines, via (50), an unique equivalence class of affine mappings from \mathbb{R}^m to \mathbb{R}^d . Consequently, $[g_c]_{\sim} = [f_s]_{\sim}$, since these classes contain n -rank affine maps that hold the same $n + 1$ linear equations. Thus, by the definition of distortion measures

$$\mathcal{D}(dg_c) = \mathcal{D}(g_c) = \mathcal{D}(f_s) = \mathcal{D}(df_s). \tag{91}$$

According to (90), the proof is completed by taking a weighted sum of the left-handed and right-handed sides of (91), over all simplices $s \in \mathcal{S}$. \square

Broadly speaking, subsequent refinements of a simplicial map f form a series of map representations in increasing resolutions, i.e., f is represented with respect to an increasing number of simplices with diminishing sizes. In this context, Theorem 9.1 establishes an important property of a multi-resolution invariance of volume-weighted

distortion energies. In order to formulate a more general multi-resolution invariance, we define an equivalence between simplicial complexes and simplicial mappings:

Definition 9.3 Two simplicial complexes $\mathcal{C}_1 = (\mathcal{C}_1, \mathcal{U}_1, \Psi_1)$ and $\mathcal{C}_2 = (\mathcal{C}_2, \mathcal{U}_2, \Psi_2)$ are *equivalent* if both are refinements of some simplicial complex $\mathcal{S} = (\mathcal{S}, \mathcal{V}, \Psi)$. Similarly, simplicial maps f_1 and f_2 of \mathcal{C}_1 and \mathcal{C}_2 are called *equivalent* if they are refinements of some simplicial map g of \mathcal{S} .

Theorem 9.1 leads to the following immediate conclusion:

Proposition 9.1 The values of volume-weighted distortion energies $E_{\mathcal{D}}(f)$ and $E_{\mathcal{D}}(g)$ are equal for two equivalent simplicial maps f and g .

Practically, 3D models are often represented by a coarse template mesh, equipped with an iterative procedure for mesh refining. This procedure is often called a *subdivision scheme*. Numerous subdivision schemes for polygonal meshes have been developed in geometric modeling for obtaining an efficient multi-resolution representations of models.

Among many others, modeling of subdivision surfaces includes tessellation schemes, also called *linear subdivision*

schemes of triangular meshes. These schemes produce equivalent simplicial complexes that induce equivalent simplicial maps. For instance, in surface mapping tasks, each triangle τ of a triangulated surface can be subdivided into four smaller triangles through the edge midpoints of τ . A recent study of [81] noticed that, in these scenarios, geometric distortions are preserved during the transition between multiple resolutions. Theorem 9.1 and Proposition 9.1 generalize and prove formally these findings for arbitrary linear subdivision schemes, delineated over n -dimensional simplicial complexes. Unfortunately, we cannot employ Theorem 9.1 for more general subdivision methods (*non-linear*), such as Loop [70] or triangulated Catmull-Clark schemes [106] (Fig. 20), where shapes of subdivided surfaces differ from the shape of the original surface. Nevertheless, as shown in Fig. 20, subdividing maps between triangular meshes, via non-linear schemes, either decreases values of distortion measures, or preserves these values in a narrow range within their original values. The qualitative explanation for this phenomena is addressed below.

Common subdivision schemes are designed for attaining smooth surfaces with a high triangulation quality. Usually, mesh triangles become nearly regular after a number of subdivision steps. This, in turn, leads to low conformal distortions, due to the almost-identical shapes of obtained source and target triangles. In particular, if each affine component of a simplicial map f is a similarity transformation, then $\sigma_1(df_s) = \sigma_2(df_s) = \dots = \sigma_n(df_s)$ for all $s \in \mathcal{S}$. Therefore, in this case, conformal distortions (26) (27), (29) and (30) attain their global minimum.

Isometry-based distortions are bounded by the ratio of relative sizes of source and target domains. If these sizes are the same, then, similarly to the conformal case, a subdivision process is capable of decreasing isometric distortions because it produces the same number of the source and target simplices at each iteration. Otherwise, isometric distortions should approach some lower bound. This bound depends on the scaling factors and total number of simplices, presented in the lowest resolution.

Clearly, using subdivision schemes for minimizing distortions is not a practical approach because subdivision process leads to an exponential growth in the number of simplices.

Nevertheless, according to the above property of simplex subdivisions, one can first minimize distortion for simplicial mapping of a coarse triangulation and then subdivide source and target domains for obtaining a low distortion map in a higher resolution without degrading the results. Moreover, as shown by Fig. 20, values of most common distortions are significantly reduced after few iterations of triangle subdivisions.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10851-021-01038-y>.

Acknowledgements This research has been supported by the Ollendorff Minerva Center and by the Office of the Chief Scientist, Israeli Ministry of Economy. AN was supported at various stages of the research by Ollendorff and Jacobs-Qualcomm Fellowships.

Appendix A Computational Aspects

We are now in the position to deliver some detail on the implementation. We explicitly write down in full matrix form the main expressions that arise in Sect. 4.

Appendix A.1 Jacobian of Simplicial Map

First, assume a general scenario (45) in which dimensions m and d of the embedding spaces may differ from the simplex dimension.

Let (v_1, \dots, v_{n+1}) be an oriented n -simplex $s \in \mathcal{S}$. To simplify our notations, we denote by v_i an element of the vertex set (*vertex index*) and by \mathbf{v}_i we denote the source coordinates of that vertex. That is, if \mathbf{y} is the source coordinate vector, then $\mathbf{v}_i \triangleq \mathbf{y}_{v_i} \in \mathbb{R}^m$. In particular, by \mathbf{v}_{ij} we denote the j -th coordinate of the point \mathbf{y}_{v_i} . We employ similar notations for the target coordinates \mathbf{x} , i.e., $f_s(\mathbf{v}_i) = f_s(v_i) = \mathbf{x}_{v_i}$, and $f(v_i)_j$ denotes the j -th coordinate of the point $\mathbf{x}_{v_i} \in \mathbb{R}^d$.

Assuming that vertex coordinates are column vector, one can then represent simplex s as a matrix $M_s \in \mathbb{R}^{m \times (n+1)}$

$$M_s \triangleq \begin{pmatrix} \mathbf{v}_{11} & \mathbf{v}_{21} & \dots & \mathbf{v}_{(n+1)1} \\ \mathbf{v}_{12} & \mathbf{v}_{22} & \dots & \mathbf{v}_{(n+1)2} \\ \vdots & \vdots & \dots & \vdots \\ \mathbf{v}_{1m} & \mathbf{v}_{2m} & \dots & \mathbf{v}_{(n+1)m} \end{pmatrix},$$

where the order of vertices (columns) reflects simplex orientation. Then, the image of s under f_s is represented by

$$M_{s'} \triangleq \begin{pmatrix} f_s(v_1)_1 & f_s(v_2)_1 & \dots & f_s(v_{n+1})_1 \\ f_s(v_1)_2 & f_s(v_2)_2 & \dots & f_s(v_{n+1})_2 \\ \vdots & \vdots & \dots & \vdots \\ f_s(v_1)_d & f_s(v_2)_d & \dots & f_s(v_{n+1})_d \end{pmatrix}. \tag{92}$$

Using hat functions (48), the expression for the map becomes

$$f_s(\mathbf{r}) = M_{s'} \mathbf{h}_s(\mathbf{r}), \quad \mathbf{r} \in \text{conv}(s),$$

where

$$\mathbf{h}_s(\mathbf{r}) \triangleq [h_{v_1}(\mathbf{r}) \ h_{v_2}(\mathbf{r}) \ \dots \ h_{v_{n+1}}(\mathbf{r})]^\top \in \mathbb{R}^{(n+1) \times 1}. \tag{93}$$

Let ∇h_{v_i} be the gradient of h_{v_i} , computed with respect to $\mathbf{r} \in \text{conv}(s)$, then ∇h_{v_i} is constant in $\text{conv}(s)$ and according to (51)

$$dh_{v_i} = -\frac{\boldsymbol{\eta}_i}{n \text{Vol}(s)} \in \mathbb{R}^{m \times 1},$$

where $\boldsymbol{\eta}_i$ is a normal vector defined in (51). Consequently, the Jacobian matrix of $\mathbf{h}_s(\mathbf{r})$, denoted by $d\mathbf{h}_s$, is given by

$$d\mathbf{h}_s = [dh_{v_1}(\mathbf{v}) \ dh_{v_2}(\mathbf{v}) \ \cdots \ dh_{v_{n+1}}(\mathbf{v})]^\top \in \mathbb{R}^{(n+1) \times m}, \quad (94)$$

and the Jacobian of f_s is

$$df_s = M_{s'} d\mathbf{h}_s, \quad df_s \in \mathbb{R}^{d \times m}. \quad (95)$$

Note that $\text{rank}(df_s) \leq n$. Therefore, df_s has at most n non-zero singular values $\sigma_1(df_s), \dots, \sigma_n(df_s)$. These singular values can be used for computing n -dimensional distortions of df_s . Equivalently, we can repeat the above computations for the canonical representation \tilde{f}_s of the map for obtaining the canonical form of the Jacobian, $d\tilde{f}_s \in \mathbb{R}^{n \times n}$. Jacobian df_s and its canonical form have the same n singular values,

$$\sigma_1(d\tilde{f}_s) = \sigma_1(df_s), \dots, \sigma_n(d\tilde{f}_s) = \sigma_n(df_s),$$

so that $\mathcal{D}(\tilde{f}_s) = \mathcal{D}(\sigma_1(df_s), \dots, \sigma_n(df_s))$ for any distortion \mathcal{D} .

Appendix A.2 Distortion Energy Gradient

Assume the scenario (43) of equal dimensions, $m = n = d$. Denote by $J \in \mathbb{R}^{n \times n}$ the Jacobian matrix df_s of a simplicial map f on a simplex s , and denote by $U \text{diag}(\sigma) V^\top$ the SVD of J . According to [89] and [40] the derivation of a distortion \mathcal{D} , computed with respect to J , is given by

$$\nabla_J \mathcal{D}(J) = U \text{diag}(\nabla_\sigma \mathcal{D}(\sigma)) V^\top. \quad (96)$$

If $M_{s'} \in \mathbb{R}^{n \times (n+1)}$ is the matrix (92) of target coordinates of s , then the derivation of $E_{\mathcal{D}}$ with respect to $M_{s'}$ can be written as:

$$\nabla_{M_{s'}} E_{\mathcal{D}}(M_{s'}) = \text{Vol}(s) d\mathbf{h} \nabla_{J_s} \mathcal{D}(J_s) \in \mathbb{R}^{(n+1) \times n}, \quad (97)$$

where J_s denotes the Jacobian $df_s[\mathbf{x}]$ and $d\mathbf{h}$ is defined according to (94). The gradient $\nabla E_{\mathbf{x}}(\mathbf{x}) \in \mathbb{R}^{n|\mathcal{V}|}$ of E is computed with respect to the column vector \mathbf{x} by laying elements of (97) into their positions in the column vector,

$$\nabla_{M_{s'}} E(M_{s'}) \mapsto \nabla E_{\mathbf{x}}(M_{s'}) \in \mathbb{R}^{n|\mathcal{V}|},$$

and then $\nabla E_{\mathbf{x}}(M_{s'})$ are summed for every simplex $s \in \mathcal{S}$.

Note that (96) is also valid for the signed SVD representation of distortion measures, mentioned in Remark 2.1. See the supplemental material for a more general computation of $\nabla_{\mathbf{x}} E$.

References

- Aharon, I., Chen, R., Zorin, D., Weber, O.: Bounded distortion tetrahedral metric interpolation. *ACM Trans. Graph.* **38**(6), 182 (2019)
- Aigerman, N., Kovalsky, S.Z., Lipman, Y.: Spherical orbifold tutte embeddings. *ACM Trans. Graph.* **36**(4), 90 (2017)
- Aigerman, N., Lipman, Y.: Injective and bounded distortion mappings in 3D. *ACM Trans. Graph.* **32**(4), 1061–10614 (2013)
- Aigerman, N., Lipman, Y.: Orbifold tutte embeddings. *ACM Trans. Graph.* **34**(6), 190–1 (2015)
- Aigerman, N., Lipman, Y.: Hyperbolic orbifold tutte embeddings. *ACM Trans. Graph.* **35**(6), 217–1 (2016)
- Aigerman, N., Poranne, R., Lipman, Y.: Lifted bijections for low distortion surface mappings. *ACM Trans. Graph.* **33**(4), 69 (2014)
- Aigerman, N., Poranne, R., Lipman, Y.: Seamless surface mappings. *ACM Trans. Graph.* **34**(4), 1–13 (2015)
- Ben-Chen, M., Gotsman, C.: Characterizing shape using conformal factors. In: 3DOR, pp. 1–8 (2008)
- Bright, A., Chien, E., Weber, O.: Harmonic global parametrization with rational holonomy. *ACM Trans. Graph.* **36**(4), 1–15 (2017)
- Bouaziz, S., Deuss, M., Schwartzburg, Y., Weise, T., Pauly, M.: Shape-up: Shaping discrete geometry with projections. *Comput. Graph. Forum* **31**, 1657–1667 (2012)
- Beck, A.: First-order methods in optimization. SIAM, Philadelphia (2017)
- Boyer, D.M., Lipman, Y., Clair, E.S., Puente, J., Patel, B.A., Funkhouser, T., Jernvall, J., Daubechies, I.: Algorithms to automatically quantify the geometric similarity of anatomical surfaces. *Proc. Natl. Acad. Sci.* **108**(45), 18221–18226 (2011)
- Bouaziz, S., Martin, S., Liu, T., Kavan, L., Pauly, M.: Projective dynamics: fusing constraint projections for fast simulation. *ACM Trans. Graph.* **33**(4), 1–11 (2014)
- Bommers, D., Zimmer, H., Kobbelt, L.: Mixed-integer quadrangulation. *ACM Trans. Graph.* **28**(3), 1–10 (2009)
- Caraman, P.: n -dimensional quasiconformal (QCF) mappings. Revised, enlarged and translated from the Roumanian by the author. (1974)
- Campen, M., Bommers, D., Kobbelt, L.: Quantized global parametrization. *ACM Trans. Graph. (TOG)* **34**(6), 1–12 (2015)
- Claici, S., Bessmeltsev, M., Schaefer, S., Solomon, J.: Isometry-aware preconditioning for mesh parameterization. In: *Computer Graphics Forum*, vol. 36, Wiley Online Library, pp. 37–47 (2017)
- Chen, R., Gotsman, C.: Generalized as-similar-as-possible warping with applications in digital photography. *Comput. Graph. Forum* **35**, 81–92 (2016)
- Choi, P.T., Lam, K.C., Lui, L.M.: Flash: Fast landmark aligned spherical harmonic parameterization for genus-0 closed brain surfaces. *SIAM J. Imag. Sci.* **8**(1), 67–94 (2015)
- Choi, G.P., Leung-Liu, Y., Gu, X., Lui, L.M.: Parallelizable global conformal parameterization of simply-connected surfaces via partial welding. *SIAM J. Imag. Sci.* **13**(3), 1049–1083 (2020)
- Chien, E., Levi, Z., Weber, O.: Bounded distortion parametrization in the space of metrics. *ACM Trans. Graph.* **35**(6), 215 (2016)
- Crane, K., Pinkall, U., Schröder, P.: Spin transformations of discrete surfaces. In: *ACM SIGGRAPH 2011 papers*, pp. 1–10 (2011)

23. Crane, K., Pinkall, U., Schröder, P.: Robust fairing via conformal curvature flow. *ACM Trans. Graph.* **32**(4), 1–10 (2013)
24. Choi, G.P., Rycroft, C.H.: Density-equalizing maps for simply connected open surfaces. *SIAM J. Imag. Sci.* **11**(2), 1134–1178 (2018)
25. Campen, M., Shen, H., Zhou, J., Zorin, D.: Seamless parametrization with arbitrary cones for arbitrary genus. *ACM Trans. Graph.* **39**(1), 1–19 (2019)
26. Chen, R., Weber, O.: GPU-accelerated locally injective shape deformation. *ACM Trans. Graph.* **36**(6), 214 (2017)
27. Du, X., Aigerman, N., Zhou, Q., Kovalsky, S.Z., Yan, Y., Kaufman, D.M., Ju, T.: Lifting simplices to find injectivity. *ACM Trans. Graph.*
28. Dorling, D., Barford, A., Newman, M.: Worldmapper: the world as you've never seen it before. *IEEE Trans. Visualiz. Comput. Graphics* **12**(5), 757–764 (2006)
29. Degener, P., Meseth, J., Klein, R.: An adaptable surface parameterization method. *IMR* **3**, 201–213 (2003)
30. Ezuz, D., Ben-Chen, M.: Deblurring and denoising of maps between shapes. In: *Computer Graphics Forum*, vol. 36, Wiley Online Library, pp. 165–174 (2017)
31. Ezuz, D., Solomon, J., Ben-Chen, M.: Reversible harmonic maps between discrete surfaces. *ACM Trans. Graph.* **38**(2), 1–12 (2019)
32. Ebke, H.-C., Schmidt, P., Campen, M., Kobbelt, L.: Interactively controlled quad remeshing of high resolution 3d models. *ACM Trans. Graph.* **35**(6), 218 (2016)
33. Floater, M.S., et al.: Parametrization and smooth approximation of surface triangulations. *Comput. Aided Geomet. Des.* **14**(3), 231–250 (1997)
34. Floater, M., Hormann, K.: Surface parameterization: a tutorial and survey (2005)
35. Freitag, L.A., Knupp, P.M.: Tetrahedral mesh improvement via optimization of the element condition number. *Int. J. Numer. Methods Eng.* **53**(6), 1377–1391 (2002)
36. Fu, X.-M., Liu, Y.: Computing inversion-free mappings by simplex assembly. *ACM Trans. Graph.* **35**(6), 216 (2016)
37. Fu, X.-M., Liu, Y., Guo, B.: Computing locally injective mappings by advanced mips. *ACM Trans. Graph.* **34**(4), 71 (2015)
38. Floater, M.: One-to-one piecewise linear mappings over triangulations. *Math. Comput.* **72**(242), 685–696 (2003)
39. Feng, Y., Wu, F., Shao, X., Wang, Y., Zhou, X.: Joint 3d face reconstruction and dense alignment with position map regression network. In: *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 534–551 (2018)
40. Giles, M.: An extended collection of matrix derivative results for forward and reverse mode automatic differentiation
41. Gastner, M.T., Newman, M.E.: Diffusion-based method for producing density-equalizing maps. *Proc. Natl. Acad. Sci.* **101**(20), 7499–7504 (2004)
42. Golla, B., Seidel, H.-P., Chen, R.: Piecewise linear mapping optimization based on the complex view. In: *Computer Graphics Forum*, vol. 37, Wiley Online Library, pp. 233–243 (2018)
43. Gu, X., Wang, Y., Chan, T.F., Thompson, P.M., Yau, S.-T.: Genus zero surface conformal mapping and its application to brain surface mapping. *IEEE Trans. Med. Imag.* **23**(8), 949–958 (2004)
44. Hefetz, E.F., Chien, E., Weber, O.: A subspace method for fast locally injective harmonic mapping. In: *Computer Graphics Forum*, vol. 38, Wiley Online Library, pp. 105–119 (2019)
45. Hormann, K., Greiner, G.: MIPS: An efficient global parameterization method. *Tech. rep, DTIC Document* (2000)
46. Herrmann, M., Herzog, R., Schmidt, S., Vidal-Núñez, J., Wachsmuth, G.: Discrete total variation with finite elements and applications to imaging. *J. Math. Imag. Vis.* **61**(4), 411–431 (2019)
47. Horn, R.A., Johnson, C.R.: Matrix analysis. corrected reprint of the 1985 original (1990)
48. Horn, R.A., Johnson, C.R.: Matrix Analysis. Cambridge University Press, Cambridge (2012)
49. Hormann, K., Lévy, B., Sheffer, A.: Mesh parameterization: Theory and practice
50. Hu, Y., Schneider, T., Wang, B., Zorin, D., Panozzo, D.: Fast tetrahedral meshing in the wild. *ACM Trans. Graph.* **39**(4), 117 (2020)
51. Hamidian, H., Zhong, Z., Fotouhi, F., Hua, J.: Surface registration with eigenvalues and eigenvectors. *IEEE Trans. Visualiz. Comput. Graph.* (2019)
52. Hu, Y., Zhou, Q., Gao, X., Jacobson, A., Zorin, D., Panozzo, D.: Tetrahedral meshing in the wild. *ACM Trans. Graph.* **37**(4), 60–1 (2018)
53. Jiang, Z., Schaefer, S., Panozzo, D.: Simplicial complex augmentation framework for bijective maps. *ACM Trans. Graph.* **36**(6), 186 (2017)
54. Jin, M., Zeng, W., Luo, F., Gu, X.: Computing Teichmüller shape space. *IEEE Trans. Visualiz. Comput. Graph.* **15**(3), 504–517 (2009)
55. Kovalsky, S.Z., Aigerman, N., Basri, R., Lipman, Y.: Controlling singular values with semidefinite programming. *ACM Trans. Graph.* **33**(4), 68 (2014)
56. Kovalsky, S.Z., Aigerman, N., Basri, R., Lipman, Y.: Large-scale bounded distortion mappings. *ACM Trans. Graph.* **34**(6), 191 (2015)
57. Kovalsky, S.Z., Galun, M., Lipman, Y.: Accelerated quadratic proxy for geometric optimization. *ACM Trans. Graph.* **35**(4), 134 (2016)
58. Kim, V.G., Lipman, Y., Chen, X., Funkhouser, T.: Möbius transformations for global intrinsic symmetry analysis. In: *Computer Graphics Forum*, vol. 29, Wiley Online Library, pp. 1689–1700 (2010)
59. Kühnel, W., Rademacher, H.-B.: Liouville's theorem in conformal geometry. *J. Mathématiques pures et appliquées* **88**(3), 251–260 (2007)
60. Kraevoy, V., Sheffer, A.: Cross-parameterization and compatible remeshing of 3d models. *ACM Trans. Graph.* **23**(3), 861–869 (2004)
61. Kraevoy, V., Sheffer, A., Gotsman, C.: Matchmaker: constructing constrained texture maps. *ACM Trans. Graph.* **22**(3), 326–333 (2003)
62. Kharevych, L., Springborn, B., Schröder, P.: Discrete conformal mappings via circle patterns. *ACM Trans. Graph.* **25**(2), 412–438 (2006)
63. Liu, T., Bouaziz, S., Kavan, L.: Towards real-time simulation of hyperelastic materials. *arXiv preprint arXiv:1604.07378* (2016)
64. Lewis, A.S.: The convex analysis of unitarily invariant matrix functions. *J. Convex Anal.* **2**(1), 173–183 (1995)
65. Li, X., Guo, X., Wang, H., He, Y., Gu, X., Qin, H.: Harmonic volumetric mapping for solid modeling applications. In: *Proceedings of the ACM*, pp. 109–120 (2007)
66. Lipman, Y.: Bounded distortion mapping spaces for triangular meshes. *ACM Trans. Graph.* **31**(4), 1081–10813 (2012)
67. Lipman, Y.: Bijective mappings of meshes with boundary and the degree in mesh processing. *SIAM J. Imag. Sci.* **7**(2), 1263–1283 (2014)
68. Lipman, Y., Kim, V.G., Funkhouser, T.A.: Simple formulas for quasiconformal plane deformations. *ACM Trans. Graph.* **31**(5), 1241–12413 (2012)
69. Li, M., Kaufman, D.M., Kim, V.G., Solomon, J., Sheffer, A.: Optcuts: joint optimization of surface cuts and parameterization. In: *SIGGRAPH Asia 2018 Technical Papers*, ACM, p. 247 (2018)
70. Loop, C.: Smooth subdivision surfaces based on triangles
71. Lipman, Y., Puente, J., Daubechies, I.: Conformal wasserstein distance: II. computational aspects and extensions. *Math. Comput.*, pp. 331–381 (2013)

72. Lévy, B., Petitjean, S., Ray, N., Maillot, J.: Least squares conformal maps for automatic texture atlas generation. In: *Acm Transactions on Graphics* (tog)
73. Liu, L., Xu, Y., Gotsman, C., Gortler, S.J.: A local/global approach to mesh parameterization. *Comput. Graph. Forum* **27**(5), 1495–1504 (2008)
74. Liu, L., Ye, C., Ni, R., Fu, X.-M.: Progressive parameterizations. *ACM Trans. Graph.* **37**(4), 41 (2018)
75. Mandad, M., Cohen-Steiner, D., Kobbelt, L., Alliez, P., Desbrun, M.: Variance-minimizing transport plans for inter-surface mapping. *ACM Trans. Graph.* **36**(4), 39 (2017)
76. Myles, A., Pietroni, N., Zorin, D.: Robust field-aligned global parametrization. *ACM Trans. Graph.* **33**(4), 1–14 (2014)
77. Myles, A., Zorin, D.: Controlled-distortion constrained global parametrization. *ACM Trans. Graph.* **32**(4), 1–14 (2013)
78. Naitsat, A., Cheng, S., Qu, X., Fan, X., Saucan, E., Zeevi, Y.Y.: Geometric approach to detecting volumetric changes in medical images. *J. Comput. Appl. Math.* **329**, 37–50 (2018)
79. Naitsat, A., Saucan, E., Zeevi, Y.Y.: Volumetric quasi-conformal mappings - quasi-conformal mappings for volume deformation with applications to geometric modeling. *Proc. VISIGRAPP 2015*, 46–57 (2015)
80. Naitsat, A., Saucan, E., Zeevi, Y.Y.: Geometric approach to estimation of volumetric distortions. In: *Proceedings of VISIGRAPP*, (2016)
81. Naitsat, A., Saucan, E., Zeevi, Y.Y.: Geometry-based distortion measures for space deformation. *Graph. Models* **100**, 12–25 (2018)
82. Naitsat, A., Zeevi, Y.Y.: Multi-resolution approach to computing locally injective maps on meshes. In: *ACM SIGGRAPH 2019 Posters*, pp. 1–2 (2019)
83. Naitsat, A., Zhu, Y., Zeevi, Y.Y.: Adaptive block coordinate descent for distortion minimization. In: *Computer Graphics Forum*, Wiley Online Library (2020)
84. Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., Guibas, L.: Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph.* **31**(4), 1–11 (2012)
85. Ouyang, W., Peng, Y., Yao, Y., Zhang, J., Deng, B.: Anderson acceleration for nonconvex admm based on douglas-rachford splitting. *Comput. Graphics Forum* **39**, 221–239 (2020)
86. Peng, Y., Deng, B., Zhang, J., Geng, F., Qin, W., Liu, L.: Anderson acceleration for geometry optimization and physics simulation. *ACM Trans. Graph.* **37**(4), 1–14 (2018)
87. Poranne, R., Tarini, M., Huber, S., Panozzo, D., Sorkine-Hornung, O.: Autocuts: simultaneous distortion and cut optimization for uv mapping. *ACM Trans. Graph.* **36**(6), 215 (2017)
88. Pach, J., Wenger, R.: Embedding planar graphs at fixed vertex locations. *Graphs Combin.* **17**(4), 717–728 (2001)
89. Rabinovich, M., Poranne, R., Panozzo, D., Sorkine-Hornung, O.: Scalable locally injective mappings. *ACM Trans. Graph.* **36**(2), 16 (2017)
90. Sorkine, O., Alexa, M.: As-rigid-as-possible surface modeling. In: *Symposium on Geometry processing*, vol. 4 (2007)
91. Saucan, E., Appleboim, E., Barak-Shimron, E., Lev, R., Zeevi, Y.Y.: Local versus global in quasi-conformal mapping for medical imaging. *J. Math. Imag. Vis.* **32**(3), 293–311 (2008)
92. Schreiner, J., Asirvatham, A., Praun, E., Hoppe, H.: Inter-surface mapping. In: *ACM SIGGRAPH 2004 Papers*, pp. 870–877 (2004)
93. Schmidt, P., Born, J., Campen, M., Kobbelt, L.: Distortion-minimizing injective maps between surfaces. *ACM Trans. Graph.* **38**(6), 1–15 (2019)
94. Sawhney, R., Crane, K.: Boundary first flattening. *ACM Trans. Graph.* **37**(1), 5 (2018)
95. Sorkine, O., Cohen-Or, D., Goldenthal, R., Lischinski, D.: Bounded-distortion piecewise mesh parameterization. In: *IEEE Visualization, 2002. VIS 2002.*, IEEE, pp. 355–362 (2002)
96. 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. *Comput. Aided Geom. Des.* **46**, 76–91 (2016)
97. Sethian, J.A.: Fast marching methods. *SIAM Rev.* **41**(2), 199–235 (1999)
98. Su, J.-P., Fu, X.-M., Liu, L.: Practical foldover-free volumetric mapping construction. In: *Computer Graphics Forum*, vol. 38, Wiley Online Library, pp. 287–297 (2019)
99. Shen, H., Jiang, Z., Zorin, D., Panozzo, D.: Progressive embedding. *ACM Trans. Graph.* **38**(4), 32–1 (2019)
100. Sheffer, A., Lévy, B., Mogilnitsky, M., Bogomyakov, A.: Abf++: fast and robust angle based flattening. *ACM Trans. Graph.* **24**(2), 311–330 (2005)
101. Solomon, J.: *Optimal transport on discrete domains*. AMS Short Course on Discrete Differential Geometry. (2018)
102. Shtengel, A., Poranne, R., Sorkine-Hornung, O., Kovalsky, S.Z., Lipman, Y.: Geometric optimization via composite majorization. *ACM Trans. Graph.* **36**(4), 38 (2017)
103. Smith, J., Schaefer, S.: Bijective parameterization with free boundaries. *ACM Trans. Graph.* **34**(4), 70 (2015)
104. Soliman, Y., Slepčev, D., Crane, K.: Optimal cone singularities for conformal flattening. *ACM Trans. Graph.* **37**(4), 1–17 (2018)
105. Springborn, B., Schröder, P., Pinkall, U.: Conformal equivalence of triangle meshes. In: *ACM SIGGRAPH 2008 papers*, pp. 1–11 (2008)
106. Stam, J.: Exact evaluation of catmull-clark subdivision surfaces at arbitrary parameter values. In: *Siggraph*, vol. 98, Citeseer, pp. 395–404 (1998)
107. Su, J.-P., Ye, C., Liu, L., Fu, X.-M.: Efficient bijective parameterizations. *ACM Trans. Graph.* **39**(4), 111–1 (2020)
108. Teran, J., Sifakis, E., Irving, G., Fedkiw, R.: Robust quasistatic finite elements and flesh simulation. In: *Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on Computer animation*, ACM, pp. 181–190 (2005)
109. Tutte, W.T.: How to draw a graph. *Proc. Lond. Math. Soc.* **3**(1), 743–767 (1963)
110. Väisälä, J.: *Lectures on n-Dimensional Quasiconformal Mappings*. Springer, Berlin (1971)
111. Vaxman, A., Müller, C., Weber, O.: Conformal mesh deformations with möbius transformations. *ACM Trans. Graph.* **34**(4), 1–11 (2015)
112. Vaxman, A., Müller, C., Weber, O.: Regular meshes from polygonal patterns. *ACM Trans. Graph.* **36**(4), 1–15 (2017)
113. Weber, O., Gotsman, C.: Controllable conformal maps for shape deformation and interpolation. In: *ACM SIGGRAPH 2010 papers*, pp. 1–11 (2010)
114. Wang, Y., Gu, X., Yau, S.-T., et al.: Volumetric harmonic map. *Commun. Inf. Syst.* **3**, 191–202 (2003)
115. Weber, O., Myles, A., Zorin, D.: Computing extremal quasiconformal maps. *Comput. Graph. Forum* **31**(5), 1679–1689 (2012)
116. Windheuser, T., Schlickewei, U., Schmidt, F.R., Cremers, D.: Geometrically consistent elastic matching of 3d shapes: A linear programming solution. In: *2011 International Conference on Computer Vision, IEEE*, pp. 2134–2141 (2011)
117. Wang, H., Yang, Y.: Descent methods for elastic body simulation on the GPU. *ACM Trans. Graph.* **35**(6), 212 (2016)
118. Weber, O., Zorin, D.: Locally injective parametrization with arbitrary fixed boundaries. *ACM Trans. Graph.* **33**(4), 1–12 (2014)
119. Xu, Y., Chen, R., Gotsman, C., Liu, L.: Embedding a triangular graph within a given boundary. *Comput. Aided Geomet. Des.* **28**(6), 349–356 (2011)
120. Yan, G., Li, W., Yang, R., Wang, H.: Inexact descent methods for elastic parameter optimization. In: *SIGGRAPH Asia 2018 Technical Papers*, ACM, p. 253 (2018)

121. Zhu, Y., Bridson, R., Kaufman, D.M.: Blended cured quasi-newton for distortion optimization. *ACM Trans. Graph.* **37**(4), 40 (2018)
122. 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*, Springer, pp. 391–408 (2009)
123. Zhang, J., Peng, Y., Ouyang, W., Deng, B.: Accelerating admm for efficient simulation and optimization. *ACM Trans. Graph.* **38**(6), 1–21 (2019)
124. Zhao, X., Su, Z., Gu, X.D., Kaufman, A., Sun, J., Gao, J., Luo, F.: Area-preservation mapping using optimal mass transport. *IEEE Trans. Visualiz. Comput. Graph.* **19**(12), 2838–2847 (2013)
125. Zhao, H., Su, K., Li, C., Zhang, B., Yang, L., Lei, N., Wang, X., Gortler, S.J., Gu, X.: Mesh parametrization driven by unit normal flow. In: *Computer Graphics Forum*, vol. 39, Wiley Online Library, pp. 34–49 (2020)
126. Zhou, J., Tu, C., Zorin, D., Campen, M.: Combinatorial construction of seamless parameter domains. *Comput. Graphics Forum* **39**, 179–190 (2020)



Prof. Zeevi is a member of the Viterbi Faculty of Electrical Engineering, Technion-Israel Institute of Technology, where he established the multidisciplinary field and research center devoted to vision and image sciences and later founded the Ollendorff Minerva Center. In 1988 he became the incumbent of the Norman and Barbara Chair in Computer Sciences. He served as the Dean of the Faculty of EE (1994-1999) and led the department through the period of its greatest growth in the field of Computer Engineering. Prof. Zeevi is the co-inventor of over 100 patents and patent applications. He is a co-founder of Cortica Ltd, a world leader in autonomous AI, and serves as its Chief Scientist. <https://webee.technion.ac.il/people/zeevi/>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Alexander Naitzat received his Ph.D. in Electrical Engineering and M.Sc. degree in Applied Mathematics from the Technion—Israel Institute of Technology. Prior to that, he obtained his B.Sc. degree in Computer Science with Mathematics at the Technion. His main research interests are Computer Vision, Geometry Processing and Computer Graphics.

<https://alexnaitzat.github.io/>



Gregory Naitzat had completed his Ph.D. at the University of Chicago, in the department of Statistics. In his research he has applied non-standard tools from abstract algebra and topology to the study of neural networks. His research was conducted under the supervision of Prof. Lek-Heng Lim. Currently his interests have shifted from academic research towards industry applications.

<https://sgregnt.github.io/>