J Math Imaging Vis (2008) 32: 293–311
DOI 10.1007/s10851-008-0101-6

Local versus Global in Quasi-Conformal Mapping for Medical Imaging

Emil Saucan · Eli Appleboim · Efrat Barak-Shimron · Ronen Lev · Yehoshua Y. Zeevi

Published online: 14 June 2008

Abstract A method and algorithm of flattening folded surfaces, for two-dimensional representation and analysis of medical images, are presented. The method is based on an application to triangular meshes of classical results of Gehring and Väisälä regarding the existence of quasi-conformal and quasi-isometric mappings.

The proposed algorithm is basically local and, therefore, suitable for extensively folded surfaces encountered in medical imaging. The theory and algorithm guarantee minimal distance, angle and area distortion. Yet, the algorithm is relatively simple, robust and computationally efficient, since it does not require computational derivatives. Both random-starting-point and curvature-based versions of the algorithm are presented.

We demonstrate the algorithm using medical data obtained from real CT images of the colon and MRI scans of the human cortex. Further applications of the algorithm, for image processing in general are also considered. The globality of this algorithm is also studied, via extreme length methods for which we develop a technique of computing straightest geodesics on polyhedral surfaces.

Keywords Surface flattening · Virtual colonoscopy · Computer aided detection · Gray-scale images · Quasiconformal mapping · Maximal dilatation · Quasi-isometry · Distortion · Conformal modulus · Quasigeodesic

E. Saucan (✉)
Electrical Engineering and Mathematics Departments, Technion, Technion City, Haifa 32000, Israel
e-mail: semil@ee.technion.ac.il

E. Appleboim · E. Barak-Shimron · Y.Y. Zeevi
Department of Electrical Engineering, Technion, Technion City, Haifa 32000, Israel

E. Appleboim
e-mail: eliap@ee.technion.ac.il

E. Barak-Shimron
e-mail: efratbarak@hotmail.com

Y.Y. Zeevi
e-mail: zeevi@ee.technion.ac.il

R. Lev
Daz 3D, Draper, UT, USA
e-mail: jupi32000@gmail.com

1 Introduction

2D representation by flattening of 3D object scans is a fundamental step, required in medical volumetric imaging. For example, it is often required to analyze three-dimensional fMRI scans of the brain cortex via their two-dimensional projections. In the latter case it is possible, for example, to better observe and follow up developments of neural activity within the folds. Flattening of 3D-scans is in particular important in CT virtual colonoscopy; a non-invasive, rapidly advancing, imaging procedure capable of determining the presence of colon pathologies such as small polyps [14]. However, because of the extensive folding of the colon, rendering of the 3D data for detection of pathologies requires the implementation of cylindrical or planar map projections [15]. In order to map such data in a meaningful manner, so that diagnosis will be accurate, it is essential that the geometric distortion, in terms of angles and lengths, caused by the representation, will be minimal. Due to these biological or biomedical applications, this problem has attracted a great deal of attention in the last few years.

1.1 Related Studies

We briefly review some of the methods, previously proposed for dealing with the flattening problem.

1.1.1 Variational Methods

Haker et al. [14, 15] introduced the use of a variational method for conformal flattening of CT/MRI 3-D scans of the brain or colon for the purpose of medical diagnosis. The method is based essentially on solving the Dirichlet problem for the Laplace-Beltrami operator $\triangle u = 0$ on a given surface Σ, with boundary conditions on $\partial\Sigma$. A solution to this problem is a harmonic (thus conformal) map from the surface to the (complex) plane. The solution suggested in [14] and [15] is a *PL* (piecewise linear) approximation of the smooth solution, achieved by solving a proper system of linear equations. Further development of this method, based upon the existence of a area-minimizing diffeomorphism between surfaces, is given in [2]. Another approach based upon minimizers of certain energy functions is developed in [8].

1.1.2 Circle Packing

Hurdal et al. [19] attempted to obtain conformal maps via circle packing. This relies on the ability to approximate conformal structure on surfaces by circle packing. The authors use this method for surfaces generated from MRI brain images and conformally map them to the three possible models of geometry in dimension 2 (i.e. the 2-sphere, the Euclidean plane and the Hyperbolic plane). Further developments of the method followed, e.g. [21]. Yet, the method is applicable for surfaces which are topologically equivalent to a disk whereas the cortical surface is not. The manner in which the works above cope with this problem is to divide the cortical surface into (at least two) closed disks and apply the flattening algorithm for each one of them. Then one either has to restrict himself to small, simply connected regions of the cortical surface, of low curvature, that render small dilatation [18, 32] or, maps the whole surface, save a small neighbourhood of a certain point. However, this is mathematically equivalent to the conformal mapping of a *punctured sphere*. Unfortunately, if one attempts to approximate an actual punctured sphere by decreasing the unmapped region, the metric distortion on the punctured surface will increase dramatically. Therefore, the dilatation of the obtained mapping will, theoretically, tend to infinity, and practically, will become excessively large, rendering the result unpractical.

An additional problem arises due to the necessary assumption that the surface triangulation is homogeneous in the sense that its triangles do not deviate to much from being equilateral. Such triangulations are seldom attainable.

1.1.3 Conformal Modulus

In the sequel, [20], of their previous work, Hurdal and Stephenson employ *conformal modulus*, computed by *extremal length* methods, mainly to extract local features in the mapping of the human cortical surface.

1.1.4 Holomorphic 1*-Forms*

Gu et al. [11–13] used holomorphic 1-forms to compute global conformal structure of a smooth surface of arbitrary genus given as a triangulated mesh. Holomorphic 1-forms are differential forms (differential operators) on smooth manifolds, which among other things can depict conformal structures. The actual computation is done via computing homology/co-homology bases for the first homology/co-homology groups of the surface, H_1, H^1 respectively. This method indeed yields a global conformal structure hence, a conformal parameterization for the surface however, computing homology basis is extremely time consuming.

1.1.5 Angle Methods

In [31] Sheffer et al. parameterize surfaces via an angle-based method, in a way that minimizes angle distortion in the process of flattening. However, the surfaces are assumed to be approximated by cone surfaces, i.e. surfaces that are composed from cone-like neighborhoods. A further development of this method that eliminates this restriction is introduced in [37].

It should be noted that in all these methods the outcome is not a conformal map but rather a quasi-conformal map. This fact implies that all the geometric measures are preserved up to some bounded distortion. Yet, sometimes the resultant distortion may render the data to become unacceptable. This is due to the fact that, in general, a surface is only locally conformal to the plane. Therefore, if one tries to flatten a surface in some manner of globality, the best result is obtained by means of a quasi-conformal flattening. Given this fact, is desirable to obtain a precise bound on the distortion caused by the flattening. However, none of these methods is capable of yielding such a bound.

The main algorithm presented in this paper proposes yet another solution to this problem. The proposed method relies on theoretical results obtained by Gehring and Väisälä [10], who studied the existence of quasi-conformal maps between Riemannian manifolds. On a practical level, its motivation resides in the fact that images are, in fact, projections on screens, such as a photographic medium, computer monitor or retina. The basic advantages of the proposed method resides in its simplicity, in setting, implementation and its speed. Additionally, it is possible to guarantee not to have distortion above a predetermined bound, which can

be as small as desired, with respect to the amount of localization one is willing to pay (and, in the case of triangulated surfaces, to the quality of the given mesh). In fact, the proposed method is—to the best of our knowledge—the only algorithm capable of computing both length distortion and angle dilatation. The algorithm introduced in this paper is most suitable for cases where the surface is complex (high and non-constant curvature) such as cortical surface or colon wrapping, or of large genus, such as skeleta, proteins, etc. Moreover, since together with the angular dilatation, both length and area distortions are readily computable, the algorithm is ideal for applications in oncology, where such measurements are extremely important.

The proposed method is basically local. In order to assess the extent at which this method can be made global we use the conformal modulus as measure. By this we extend our previous work on the local aspects of quasi-conformal mapping of folded surfaces (see, e.g. [3]).

The results of Gehring and Väisälä were adapted to triangular meshes (i.e. piecewise flat simplicial 2-manifolds). This restriction is due to our desire to devise and test methods applicable to topological/geometric data as close as possible to the "raw" data as provided by the scanner or other medical imaging device. However, if higher regularity of the surfaces under investigation is supposed and/or desirable (e.g. for better graphics quality), the algorithm can readily be modified, using standard methods (such as B-splines), so that differential notions appearing in the original work of Gehring and Väisälä can be computed.

We test this approach on CT-scans of the human colon. This requires the computation of geodesics on the colon surface which is a topological cylinder. In this paper we propose a modified version of an algorithm due to Polthier and Schmies [26] for finding *straightest geodesics* on triangular meshes.

The paper is organized as follows: In Sect. 2 we describe the necessary background material: quasi-conformal and quasi-isometric mappings, admissible surfaces, moduli of quadrilaterals and rings and the extremal length method; as well as some basic facts on quasi-geodesics on polyhedral surfaces. Section 3 is dedicated to the presentation of our main algorithms: the global one, based on conformal moduli, and where we apply a surface dissection method using *PL*-quasi-geodesics; and the local one stemming from the seminal work of Gehring and Väisälä. The subsequent section is dedicated to the analysis of the experimental results obtained by both of our main algorithms. Finally, in Sect. 5 we summarize the methods and results presented and consider a number of further developments and improvements.

2 Theoretical Background

2.1 Quasi-Conformal Mappings

We introduce below the main definitions and basic facts regarding quasi-conformal mappings in $\mathbb{R}^n$. In doing so, we are motivated by the fact that, in recent years it has become common amongst the imaging community, to consider images as higher dimensional manifolds embedded in $\mathbb{R}^n$ or $\mathbb{R}^n$ (see, e.g., [5, 16, 22, 30]). Therefore, even if somewhat technical, we bring the general definitions, as a general frame for further implementations, and do not restrict ourselves to the classical case $n = 2$, even if at this stage, our applications are limited to the 2-dimensional case. This choice is further motivated by a view towards future applications in proper (i.e. 3-dimensional) tomography.

Definition 2.1 Let $D \subseteq \mathbb{R}^n$ be a domain; $n \geq 2$ and let $f : D \to f(D)$ be a homeomorphism. f is called *quasi-conformal* iff

(i) f belongs to $W^{1,n}_{loc}(D)$, and
(ii) there exists $K \geq 1$ such that:

$$|f'(x)|^n \leq K J_f(x) \quad \text{a.e.;} \tag{2.1}$$

where $f'(x)$ denotes the formal derivative of f at x, $|f'(x)| = \sup_{|h|=1} |f'(x)h|$, and where $J_f(x) = \det f'(x)$. The smallest number K that satisfies (2.6) is called the *outer dilatation*, $K_O(f)$, of f.

One can extend the definitions above to Riemannian manifolds as follows: let M^n, N^n be to oriented, connected $\mathcal{C}^\infty$ Riemannian n-manifolds, $n \geq 2$, and let $f : M^n \to N^n$ be a continuous function. One can define the formal derivative of f by using coordinate charts (see, e.g. [28]).

If f is quasi-conformal, then there exists $K' \geq 1$ such that the following inequality holds a.e. in M^n:

$$J_f(x) \leq K' \cdot \inf_{|h|=1} |T_x f h|^n. \tag{2.2}$$

By analogy with the outer dilatation we have the following definition:

Definition 2.2 The smallest number K' that satisfies inequality (2.2) is called the *inner dilation* $K_I(f)$ of f. The *maximal dilatation* of f is defined as $\max(K_O(f), K_I(f))$. A mapping f is called K*-quasi-conformal* iff there exists $K \geq 1$ such that $K(f) \leq K$.

The dilatations are simultaneously finite or infinite. Indeed, the following inequalities hold:

$$K_I(f) \leq K_O^{n-1}(f) \quad \text{and} \quad K_O(f) \leq K_I^{n-1}(f).$$

Remark 2.3 In the classical case $n = 2$, $K_O = K_I = K$.

2.1.1 Quasi-Isometries and Projections

Definition 2.4 Let $D \subset \mathbb{R}^n$ be a domain. A homeomorphism $f : D \to \mathbb{R}^n$ is called a *quasi-isometry* (or a *bi-lipschitz mapping*), if there exists $1 \le C < \infty$ such that

$$\frac{1}{C}|p_1 - p_2| \le |f(p_1) - f(p_2)| \le C|p_1 - p_2|, \quad \text{for all } p_1, p_2 \in D; \tag{2.3}$$

where "$|\cdot|$" denotes the standard (Euclidean) metric on $\mathbb{R}^n$. $C(f) = \min\{C \mid f \text{ is a quasi-isometry}\}$ is called the *minimal distortion* of f (in D).

Remark 2.5 Evidently, the above definition readily extends to general metric spaces. In particular, for the case of hypersurface embedded in $\mathbb{R}^n$, distances are the induced intrinsic distances on the surface.

Definition 2.6 Let D be a domain in $\mathbb{R}^n$ and let $f : D \to \mathbb{R}^n$ be a mapping. f is called a *local C-quasi-isometry* (in D) iff for every $C' > C$ and for any $x \in D$, there exists a neighbourhood $U_x \subset D$ of x, such that $f|_{U_x}$ is a quasi-isometry with maximal dilatation C'.

Local quasi-isometries have the following important physical interpretation. Let $C(D) = \sup\{C \mid \text{any } f \text{ local C-quasi-isometry is injective in D}\}$. Then, if D is regarded as an elastic body and f as deformation of D under the influence of a force field, then $C(D)$ represents the critical strain, i.e. the strain under which D collapses onto itself.

If f is a quasi-isometry, then

$$K(f) \le C(f)^{2(n-1)}. \tag{2.4}$$

It follows that any quasi-isometry is a quasi-conformal mapping (while, evidently, not every quasi-conformal mapping is a quasi-isometry).

Definition 2.7 Let $S \subset \mathbb{R}^n$ be a connected, $(n-1)$-dimensional set. S is called *admissible* (or an *admissible hypersurface*) iff for any $p \in S$, and for any $\varepsilon > 0$, there exists a neighbourhood $U_p \subset \mathbb{R}^n$ of p, and there exists a quasi-isometry $i_p : U_p \to \mathbb{R}^n$ such that $i_p(S \cap U_p) = D_p \subset \mathbb{R}^{n-1}$, where D_p is a domain, and such that $C(i_p)$ satisfies:

(i) $\sup_{p \in S} C(i_p) < \infty$; and

(ii) $\operatorname{ess\,sup}_{p \in S} C(i_p) \le 1 + \varepsilon$.

Let $S \subset \mathbb{R}^n$ be a set homeomorphic to an $(n-1)$-dimensional domain, $\vec{n}$ be a fixed unitary vector, and $p \in S$, such that there exists a neighbourhood $V \subset S$, such that $V \simeq D^{n-1}$, i.e. V is homeomorphic to D^{n-1}, where

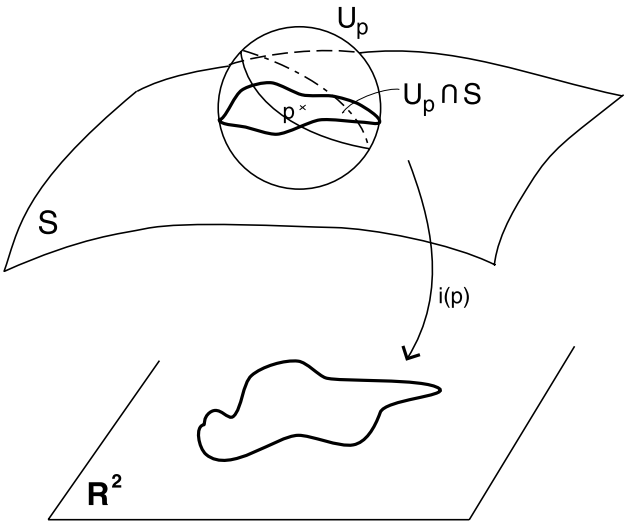

Fig. 1 An admissible surface

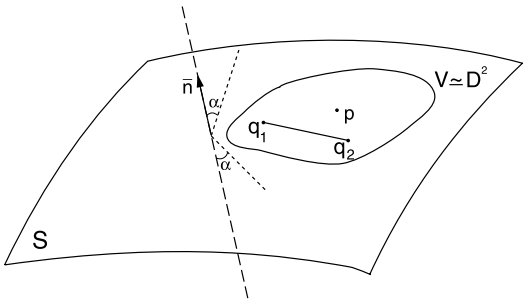

Fig. 2 The geometric condition

$D^{n-1} = \{x \in \mathbb{R}^{n-1} \mid \|x\| \le 1\}$. Moreover, suppose that for any $q_1, q_2 \in S$, the acute angle $\measuredangle(q_1 q_2, \vec{n}) \ge \alpha$. We refer to the last condition as *the Geometric Condition* or *Gehring Condition* (see [10]).

Then, for any $x \in V$ there exists a unique representation of the following form:

$$x = q_x + u\vec{n};$$

where q_x lies on the hyperplane $\Pi_{p,\vec{n}}$ through p, orthogonal to $\vec{n}$, and where $u \in \mathbb{R}$.

Define:

$$Pr(x) = q_x.$$

By [6], Corollary 1, p. 338, we have that for any $p_1, p_2 \in S$ and any $a \in \mathbb{R}_+$ the following inequalities hold:

$$\frac{a}{A}|p_1 - p_2| \le |Pr(p_1) - Pr(p_2)| \le A|p_1 - p_2|;$$

where

$$A = \frac{1}{2}[(a \csc\alpha)^2 + 2a + 1]^2 + \frac{1}{2}[(a \csc\alpha)^2 - 2a + 1]^2.$$

Fig. 3 Quadrilateral modulus

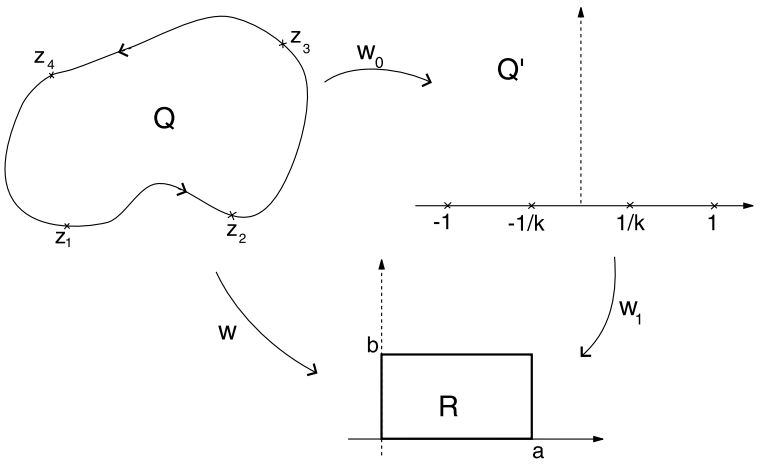

In particular, for $a = 1$ we get that

$$C(Pr) \leq \cot\alpha + 1 \tag{2.5}$$

and

$$K(Pr) \leq \left(\left(\frac{1}{2}(\cot\alpha)^2 + 4\right)^{\frac{1}{2}} + \frac{1}{2}\cot\alpha\right)^n \leq (\cot\alpha + 1)^n. \tag{2.6}$$

From the above discussion we conclude that $S \subset \mathbb{R}^3$ is an admissible hypersurface if for any $p \in S$ there exists $\vec{n}_p$, such that for any $\varepsilon > 0$ there exists $U_p \simeq D^2$, such that for any $q_1, q_2 \in U_p$ the acute angle $\measuredangle(q_1q_2, \vec{n}_p) \geq \alpha$, where

(i) $\inf_{p\in S} \alpha_p > 0$; and

(ii) $\operatorname{ess\,inf}_{p\in S} \alpha_p \geq \frac{\pi}{2} - \varepsilon$.

Example 2.8 Any hypersurface in $S \in \mathbb{R}^3$, that admits a well-defined continuous turning tangent plane at any point $p \in S$, is admissible.

2.2 Conformal Modulus and Extremal Length

2.2.1 The Conformal Modulus

Definition 2.9 A *quadrilateral* is a simply connected planar domain Q together with four points $z_1, z_2, z_3, z_4 \in \partial Q$, where the order of the points z_1, z_2, z_3, z_4 is consistent with the positive orientation on ∂Q. The points z_1, z_2, z_3, z_4 are called the *vertices* of Q and the arcs determined by z_1, z_2, z_3, z_4 on ∂Q are called the *sides* of Q: $\widehat{z_1z_2}, \widehat{z_3z_4}$ are called the *a-sides* and $\widehat{z_2z_3}, \widehat{z_4z_1}$ are called the *b-sides*.

From the Riemann Mapping Theorem (see, e.g. [24, p. 13]) it follows that any quadrilateral may be mapped on the quadrilateral $Q'(-\frac{1}{k}, -1, 1, \frac{1}{k})$, $0 < k < 1$; where Q' is the upper half-plane.

Since any three points on the boundary of Q determine the mapping completely, it follows that quadrilaterals are partitioned into equivalence classes.

By classical results in elliptic function theory it follows that

$$w_1(z) = \int_0^z \frac{d\zeta}{\sqrt{(1-\zeta^2)(1-k^2\zeta^2)}}$$

maps $Q'(-\frac{1}{k}, -1, 1, \frac{1}{k})$ onto a rectangle R. If w_0 denotes the mapping from Q onto Q', then the *canonical mapping of* Q: $w = w_1 \circ w_0$ maps the quadrilateral Q conformally onto the *canonical rectangle* $R = R(Q) = \{x + iy \,|\, 0 < x < a, 0 < y < b\}$. Therefore, any conformal class of equivalence of quadrilaterals contains rectangles. Conversely, any conformal mapping between two rectangles is a similarity transformation. Moreover, since similar rectangles obviously belong to the same equivalence class, all canonical rectangles of a given quadrilateral Q have the same ratio of the sides $a/b = M(Q)$ (hence the denomination of the a-sides and b-sides). The number $M(Q)$ is called the *conformal module* (*modulus*) of the quadrilateral Q. It follows from the discussion above that two quadrilaterals are conformally equivalent iff they have the same module.

Remark 2.10 Under circular permutations, the quadrilateral modulus displays the following behavior: $M(Q(z_1, z_2, z_3, z_4)) = M(Q(z_3, z_4, z_1, z_2)) = 1/M(Q(z_2, z_3, z_4, z_1))$.

2.2.2 Extremal Length—The Length-Area Method

The extremal length method—introduced by Beurling in 1946—represents the sought for tool for computing the

moduli of quadrilaterals (see, e.g. [24]):

$$M(Q) = \frac{\int\int |w'(z)|^2 d\sigma}{(\inf_{\beta\in\mathcal{C}_a} \int_C |w'(z)\|dz|)^2}. \tag{2.7}$$

Here $\mathcal{C}_a$ denotes the family of locally rectifiable Jordan arcs included in Q, that connect the a-sides of Q, and w denotes the canonical mapping of Q onto the canonical rectangle R.

The problem with the above estimate resides in the fact that it is given by using a certain conformal (i.e. 1-quasi-conformal mapping). Therefore a definition of the modulus of a quadrilateral that makes no appeal to quasi-conformality is desirable. Fortunately, such a definition exists and is given by the following formula (see [24]):

$$M(Q) = \inf_{\varrho\in\mathcal{P}} \frac{m_\varrho(Q)}{(\inf_{\beta\in\mathcal{C}_a} l_\varrho(\beta))^2}. \tag{2.8}$$

Here $\mathcal{P}$ denotes the family of non-negative Borel-measurable functions on Q, such that $\int_\beta \varrho(z)|dz| \geq 1$, for all $\beta \in \mathcal{C}_a$. The infimum is attained for $\varrho = |w'|$, where w is the canonical mapping of Q.

2.2.3 The Modulus of a Ring

One can also define the modulus of a *ring domain* R, i.e. of a doubly-connected planar domain. This is done in analogy to the definition of the modulus of a quadrilateral, but in this case one maps conformally R onto an *annulus* $A = \{0 \leq r_1|z| \leq r_2 \leq \infty\}$. Here we will consider only *proper* annuli, i.e. such that $r_1 \neq 0, r_2 \neq \infty$. In this non-degenerate case, the modulus of the ring R is defined as $M(R) = \log(r_2/r_1)$.

As expected, there is a close connection between the moduli of the two types of domains. Indeed, given an annulus A of radii r_1, r_2, one can cut it along the segment $[r_1, r_2]$, and conformally map it onto a quadrangle Q. The foreseen result holds, i.e. the moduli of R and Q are equal.

The extremal methods applied in the case of quadrangles are also valid for rings and rend expressions similar to (2.7) and (2.8), which, for the test function $\varrho \equiv 1$ become:

$$M(R) = 2\pi \frac{\min_c(length(c))^2}{Area(R)}; \tag{2.9}$$

and

$$M(R) = 2\pi \frac{Area(R)}{\min_\gamma(length(\gamma))^2}; \tag{2.10}$$

respectively, where α and γ are types of curves depicted in Fig. 4, i.e. γ is a locally rectifiable arc connecting the two boundary components of R, $\partial_1 R$ and $\partial_2 R$, and α is a locally rectifiable simple closed curve, separating $\partial_1 R$ and $\partial_2 R$.

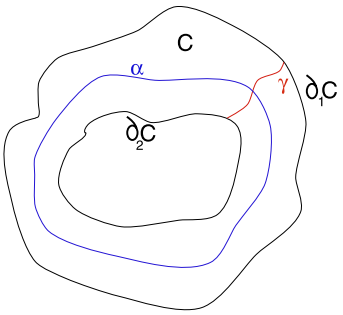

Fig. 4 Ring modulus

2.2.4 Quasi-Conformal Mappings Between Annuli

Let R, R' be annuli, and $f : R \to R'$ be a K-quasi-conformal mapping of R onto R'. Then:

$$\frac{1}{K} \leq \frac{M(R')}{M(R)} \leq K. \tag{2.11}$$

Note that, again, an analogous result holds for quadrangles.

2.3 Quasi-Geodesics on Polyhedral Surfaces

In any path-metric spaces the natural definition of geodesics is that of shortest curves. Therefore, in any combinatorial (including *PL*) context, the notion of "shortest path" designates *global shortest* path, with the obvious implementation shortcomings. Thus, one is faced with a properly classical combinatorial (and computationally hard) problem of finding paths with global shortest length.

However, in differential geometry geodesics are classically defined *locally* as *straightest* paths. This is a straightforward generalization stemming from one of the characterizations of Euclidean lines (see [17]). For smooth (differentiable) surfaces (and higher dimensional manifolds), to be a geodesic a curve must satisfy the condition: $\kappa_g \equiv 0$, where κ_g denotes the *geodesic curvature* (see [7]). It is a classical result of differential geometry that geodesics are precisely the *locally shortest paths* (see, e.g. [7], Propositions 4.6.4 and 4.6.5). Note that κ_g is an intrinsic notion, i.e it depends solely upon the inner geometry (metric) of the surface (manifold), and not upon its embedding in $\mathbb{R}^2$ ($\mathbb{R}^n$) (see, e.g. [4], Lemma 1.6.3). Also, it is important to recall that uniqueness and shortness of geodesics hold only locally, as the well known examples of the sphere and (right circular) cylinder illustrate. Moreover, even the global existence of shortest geodesics is not assured for all surfaces (manifolds), as the example of $\mathbb{R}^2 \setminus \{0\}$ with the Euclidean metric shows.

We can however try and apply the above condition $\kappa_g = 0$. Note that the partial differential equation for the

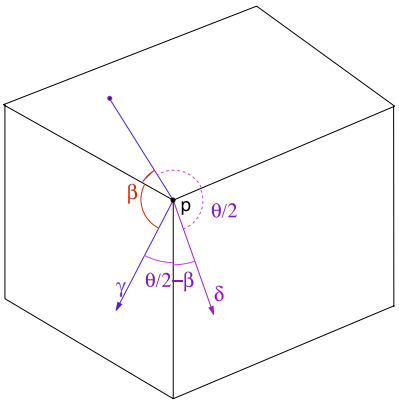

Fig. 5 Discrete geodesic curvature at a vertex of a polyhedral surface

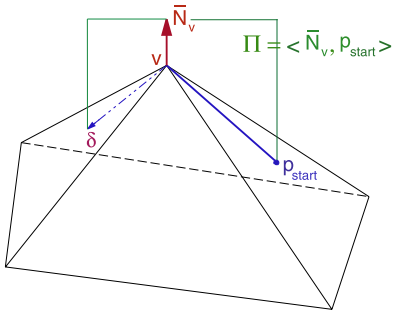

Fig. 6 Straightest geodesic through a vertex of a polyhedral surface

computation of the geodesic curvature (see, e.g. [7]), is obviously of little relevance for triangulated surfaces.[1] Instead we can apply the following formula, due to Polthier and Schmies, and an extension of previous results of the Alexandrov school (see [1, 27]), for the computation of κ_g for curves on 2-dimensional polyhedral surfaces,

$$\kappa_g = \frac{2\pi}{\theta}\left(\frac{\theta}{2} - \beta\right). \tag{2.12}$$

More precisely, let γ be a *PL* curve γ on a polyhedral surface, and let $p \in \gamma$ be a vertex of the polyhedral surface. Then the *discrete geodesic curvature* of γ at p is given by:

$$\kappa_g(\gamma, v) = \frac{2\pi}{\theta}\left(\frac{\theta}{2} - \beta\right), \tag{2.13}$$

where β represents a choice (i.e. the smallest) of the angle of γ at p—see Fig. 6 where δ represents the *discrete straightest geodesic* (at p), defined as in Fig. 7.

[1] It only serves to justify the following authoritative but, unfortunately, usually forgotten fact [4]: "There are extremely few surfaces on which geodesics (and a fortiori shortest routes) can be more or less explicitly determined."

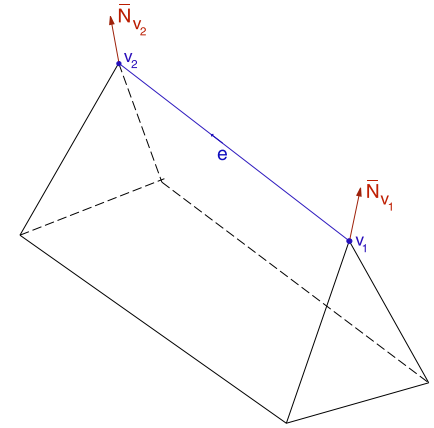

Fig. 7 Discrete geodesic curvature along an edge of a polyhedral surface

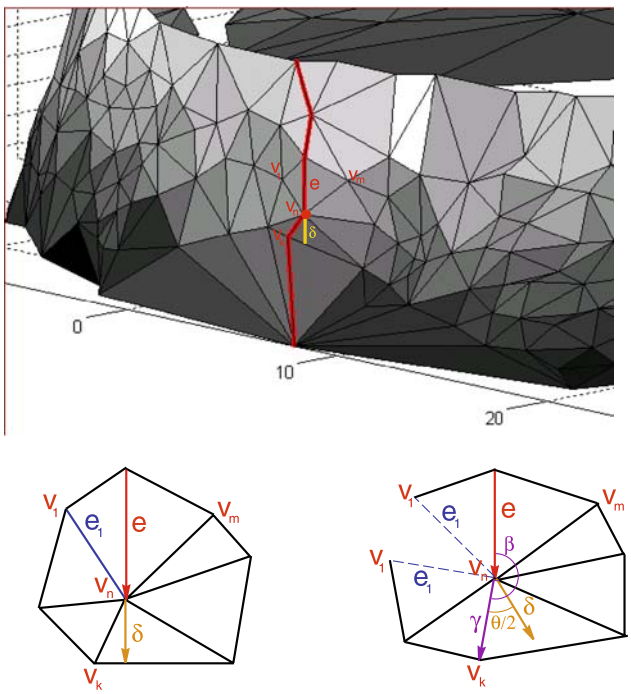

Fig. 8 Computed quasi-geodesic on the colon surface (*top*). Detail: finding the straightest geodesic using Polthier's method and determining the closest quasi-geodesic (*bottom*)

In this discrete approach to geodesic curvature, $\kappa_g(e) \equiv 0$, for any edge e. To include the geodesic curvature of the edges one can apply the following formula:

$$\kappa_g(e) = \measuredangle(\bar{N}_{v_1}, \bar{N}_{v_2}),$$

where $\bar{N}_{v_1}, \bar{N}_{v_2}$ represent the *computed* normals at the vertices that define e (see Fig. 8).

Remark 2.11 The approach described above (as quantified in (2.12) and (2.13)) originates, as we have already men-

tioned above, in the work of the Russian School. They have considered the *turn* of the curve, that is, locally, equivalent to the angle deviation of [27], and globally amounts to the *integral curvature* $\int_\gamma \kappa_g$. The following unsurprising result justifies our (and Polthier-Schmies') algorithm:

Theorem 2.12 ([1], Theorem 5.1.4) *If the turn of a* (*non-degenerate*) *curve* γ *is zero, then* γ *is a straight line segment.*

Indeed, if geodesics are characterized by the condition $\int_\gamma \kappa_g = 0$, one can define quasi-geodesic by imposing the (global) condition: $\int_\gamma \kappa_g \leq k_0$, where, in any practical implementation, k_0 can be chosen conveniently.

3 Methods and Algorithms

In this section we describe the algorithms that were developed and applied for the implementation of the theories described in previous sections. As already noted in the introduction, we have tested our algorithms on piecewise flat simplicial 2-manifolds (i.e. triangular meshes). The reason behind this restriction is our desire to devise and test methods applicable to topological/gemoetric data as close as possible to the "raw" data as provided by the scanner or other medical imaging device. Of course, if higher regularity of the surfaces under investigation is supposed and/or desirable, the algorithms can easily be adapted to this case, using standard methods (such as B-splines).

We first concentrate on the so called "global viewpoint", i.e. computation of the conformal modulus of a curved annulus.

3.1 Global Viewpoint

Before delving into the algorithm's details, we stress that the purpose of computing the conformal modulus is to obtain a first measure of the degree to which the topological cylinder under investigation (i.e. the colon surface) deviates from being a flat cylinder. This endows us, in consequence, with a (rather rough) estimate of maximal possible "globality" of the flattening algorithm.

The algorithm that was used to find quasi-geodesic paths on the triangulated approximation of the colon surface is, essentially, an implementation of a modified version of the algorithm proposed in [26] (see also the discussion on *PL*-quasi-geodesics in Sect. 2.3 above), with the minor following additional constraint that the path is allowed to run only along edges of the triangulation. Accordingly, for a given direction and a given initial vertex, the algorithm computes the straightest possible path which stretches from the given vertex in the given direction, under the above limitation that it pathes along edges. Such a constrain simplifies the algorithm significantly, since, by the considerations of Sect. 2.3, it reduces, in essence, to the computation of the angles between adjacent edges of the 1-skeleton of the triangulation. The cost of this simplification is a bounded (and controlled) divergence from the geodesic path between the two end points (hence the term quasi-geodesic). However, since we employed this method for determining "quasi-generators" on the very thin topological cylinders produced by the scanning process of the colon, the quasi-geodesics it renders are, in fact, proper geodesics (with the said proviso of running only on the edges of the triangulation).

Let us present the algorithm in some more detail: Suppose we try to determine the straightest quasi-geodesic from with start point at a vertex v_t, on one boundary component of a cylinder, and end point a vertex v_b on the other boundary component. First, we note that for the particular case of the colon surface, the algorithm is made easier by the fact that the upper boundary component of the surface (which is, al already mentioned, a topological cylinder) can be supposed to be practically parallel to the XY coordinates plane (hence the notation for the boundary points). Therefore, choosing the starting edge $v_t v_1$ can be reduced to selecting the edge of steepest descent.

Furthermore, suppose that in computing the quasi-geodesic we have already reached some vertex v_n, between v_t and v_b, adjacent to an edge $e = (v_{n-1} v_n)$ belonging to the computed quasi-geodesic. Let v_k, $k = 1, \ldots, k_0$ be the vertices adjacent to v_n, such that $v_k \neq v_{n-1}, k = 1, \ldots, k_0$. We proceed from v_n to one of its neighboring vertices v_k along an edge $v_n v_k$, where $\gamma = v_t \ldots v_n v_k$ is the quasi-geodesic extension of the $v_t v_n$, i.e. such that the path $v_p v_n v_k$ has the least deviation from the true geodesic segment $e\delta$ (see Fig. 8).

By (2.13), to achieve this, one only has to choose v_k such that $\kappa_g(v_p v_n v_k, v_n)$ is minimal. We sum up this algorithm in the pseudocode below (see also Fig. 8).

Remark 3.1 In regions where the curvature of the surface is strictly negative, the divergence of a quasi-geodesic from the geodesic is guaranteed to be bounded, whereas in regions of positive curvature, it can, theoretically, increase fast (see, e.g., [4]). However, in practice, given a fine enough resolution of the triangular mesh, the divergence remains controlled.

3.2 Local Viewpoint—Gehring-Väisälä

We present in this section an algorithm for obtaining a quasi-isometric (flat) representation of a given surface. As already mentioned above, we restrict to the case of triangular meshes. Towards this end, we have adapted the differential definition of quasi-conformality and the result of Gehring

Algorithm 1 Shortest quasi-geodesic

Input: Cylindrical triangular mesh C with its axis aligned so that the cylinder's top boundary component is almost parallel to the XY plane
Output: Set of quasi-geodesics on the mesh
$Result \leftarrow \emptyset$;
foreach *Vertex* v_t *on the upper boundary component of the cylinder* **do**
  let (v_t, v_n) be the edge of steepest descent;
  $v_p \leftarrow v_t$;
  $CurGeodesic \leftarrow \{(v_t, v_n)\}$;
  while v_n *not on bottom boundary component* **do**
    let (v_n, v_k) be the edge such that $\kappa_g(v_p v_n v_k, v_n)$ is minimal;
    $v_p \leftarrow v_n$;
    $v_n \leftarrow v_k$;
    $CurGeodesic \leftarrow CurGeodesic \cup \{(v_p, v_n)\}$;
  end
  $Result \leftarrow Result \cup CurGeodesic$
end

and Väisälä to the piecewise flat case. Again, using standard methods (e.g. B-splines, finite-element methods) one can revert to the differentiable setting and test the (almost) full analytic definition of quasi-conformality.

First, choose an angle α (according to the acceptable dilatation and distortion, cf. (2.5) and (2.6)).

Next, let N_0 stand for the normal vector to the surface at a point v_0 on the surface. There are two possibilities to choose the *starting vertex* (or *anchor vertex*) v_0, one is in a random manner and the other is based on curvature considerations. The discrete curvature measure employed is that of *angular defect*, due to its simplicity and high reliability (see [33]). We shall compare the performances of both variations of the basic algorithm in Sect. 4.2.

Let $T_i, i = 1, \ldots, m_0$ be the triangles adjacent to the vertex v_0, let $N_i, i = 1, \ldots, m_0$ denote their normals, and put $\alpha_i = \measuredangle(N_0, N_i)$. Add to the patch at v_0 all the triangles T_i such that $\alpha_i \leq \alpha$. It is easy to verify that this inequality is equivalent to the Gehring condition for the case of triangular meshes. Project the added triangles onto the computed tangent plane at v_0, too obtained the flattened patch (see Fig. 9).

If al the angles $\alpha_i = \measuredangle(N_0, N_i), i = 1, \ldots, m_0$ are larger than α (which may happen if the chosen α is very small relatively to the discrete curvature at v_0), choose any of the triangles adjacent to v_0 to be the patch at this vertex.

If the degenerate case above does not occur, we keep adding triangles T_j of the *second star* of the vertex v_0 to the patch at v_0, if the angles β_j between the normals N_j of the considered triangles, and N_0 do not exceed α (see Fig. 10). We increase the patch moving from an added triangle to its neighbors (in the second and next stars of v_0) while avoiding repetitions, till no triangles can be added (i.e. till

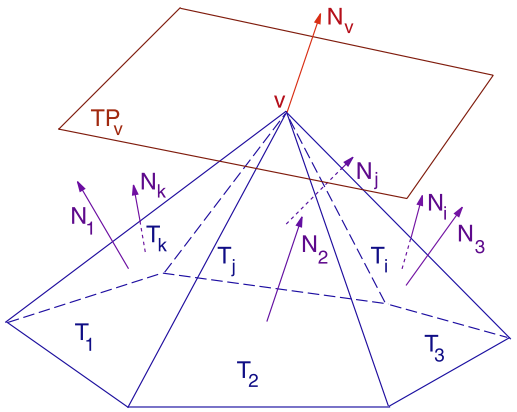

Fig. 9 Computed normal and tangent plane at a vertex of the triangular mesh

all angles are $\geq \alpha$). If by this time all triangles were added to the patch we have completed constructing the mapping. Otherwise, chose a new vertex that has not been projected yet, to be the starting triangle of a new patch (see Fig. 11). A pseudocode for this procedure can be easily written (see algorithm in the next page):

Remark 3.2 A variant of the basic algorithm above, using starting (or anchor) triangles instead of vertices, was also implemented. Of course, in the beginning, a triangle T_0, of the triangulation must be chosen. In this case, however, we project the patch at T_0 onto the plane including T_0.

After T is (trivially) projected onto itself we move to its neighbors. Suppose T_1 is a neighbor of T. We call T_1 *Gehring compatible w.r.t* T, if the maximal angle between N_{T_1} and N_T, the normal vectors to T_1 and T, respectively, does not exceed the predefined angle α. We insert T_1 to the patch of T, iff it is Gehring compatible with respect to T. As

Input: A vertex $VertexSeed$ that is the seed of the expansion
A set of vertices $CurrentVetexSet$ known to be part of the patch
$SelectedNormal$ the normal defining the patch
A predefined $MaxAngleError$ determined by the maximal acceptable dilatation/distortion
Output: Set of vertices belonging to the patch
foreach *Face* f *that includes* $VertexSeed$ *as one of its vertices* **do**
  if *Angle between normal of face* f *to* $SelectedNormal$ *is smaller than* $MaxAngleError$ **then**
    foreach *Vertex* $v \in f - CurrentVetexSet$ **do**
      if *Angle between normal of vertex* v *to* $SelectedNormal$ *is smaller than* $MaxAngleError$ **then**
        $CurrentVetexSet \leftarrow CurrentVetexSet \cup ExpandPatch\,(v, CurrentVetexSet, SelectedNormal)$;
      end
    end
  end
  return $CurrentVetexSet$
end

Function 2 `ExpandPatch(`*VertexSeed, CurrentVetexSet, SelectedNormal*`)`

Input: Triangular mesh M with estimated normals for each vertex
Output: Set of flattening maps for patches of the mesh
$Result \leftarrow \emptyset$;
$PV \leftarrow \emptyset$;
while $PV \neq V$ **do**
  select vertex $v \in V - PV$;
  let n_v be the vertex normal at vertex v;
  $CPV \leftarrow$ExpandPatch$(v, \{v\}, n_v)$;
  $PV \leftarrow PV \cup CPV$;
  $CurMap \leftarrow \emptyset$;
  foreach *vertex* k *in* CPV **do**
    $uv_k \leftarrow$projection of k according to normal n_v;
    insert mapping (k, uv_k) to $CurMap$;
  end
  insert $CurMap$ to $Result$
end

before, we keep adding triangles to the patch of T moving from an added triangle to its neighbors while avoiding repetitions, till no triangles can be added. Again, if by this time all triangles were added to the patch we have completed constructing the mapping. Otherwise, chose a new triangle that has not been included in a previous patch, to be the starting triangle of a new patch.

Remark 3.3 One should keep in mind that the above given algorithm, as for any other flattening method, is local. Indeed, in a sense the proposed algorithm gives a measure of "globality" of this intrinsically local process.

Remark 3.4 Our algorithm is best suited for highly folded surfaces, because of its intrinsic locality, on the one hand, and computational simplicity, on the other. However, on "quasi-developable" surfaces (i.e. surfaces that are almost cylindrical or conical) the algorithm behaves similar to other algorithms, with practically identical results.

4 Experimental Results

In this section we present experimental results that were obtained by implementing both the global algorithms for finding shortest quasi-geodesics and subsequent computation of the conformal modulus, and the local algorithm of flattening sampled surfaces. The local algorithm was tested on both synthetic as well as real data images of the colon and brain surfaces. The global algorithm was tested on the colon images only. All brain images were obtained from a brain MRI-imaging simulator, BrainWeb, developed at McGill university (www.bic.mni.mcgill.ca/brainweb/). All colon images are taken from colon CT-scan done in Rambam medical center, Haifa, curtesy of Dr. Doron Fisher MD.

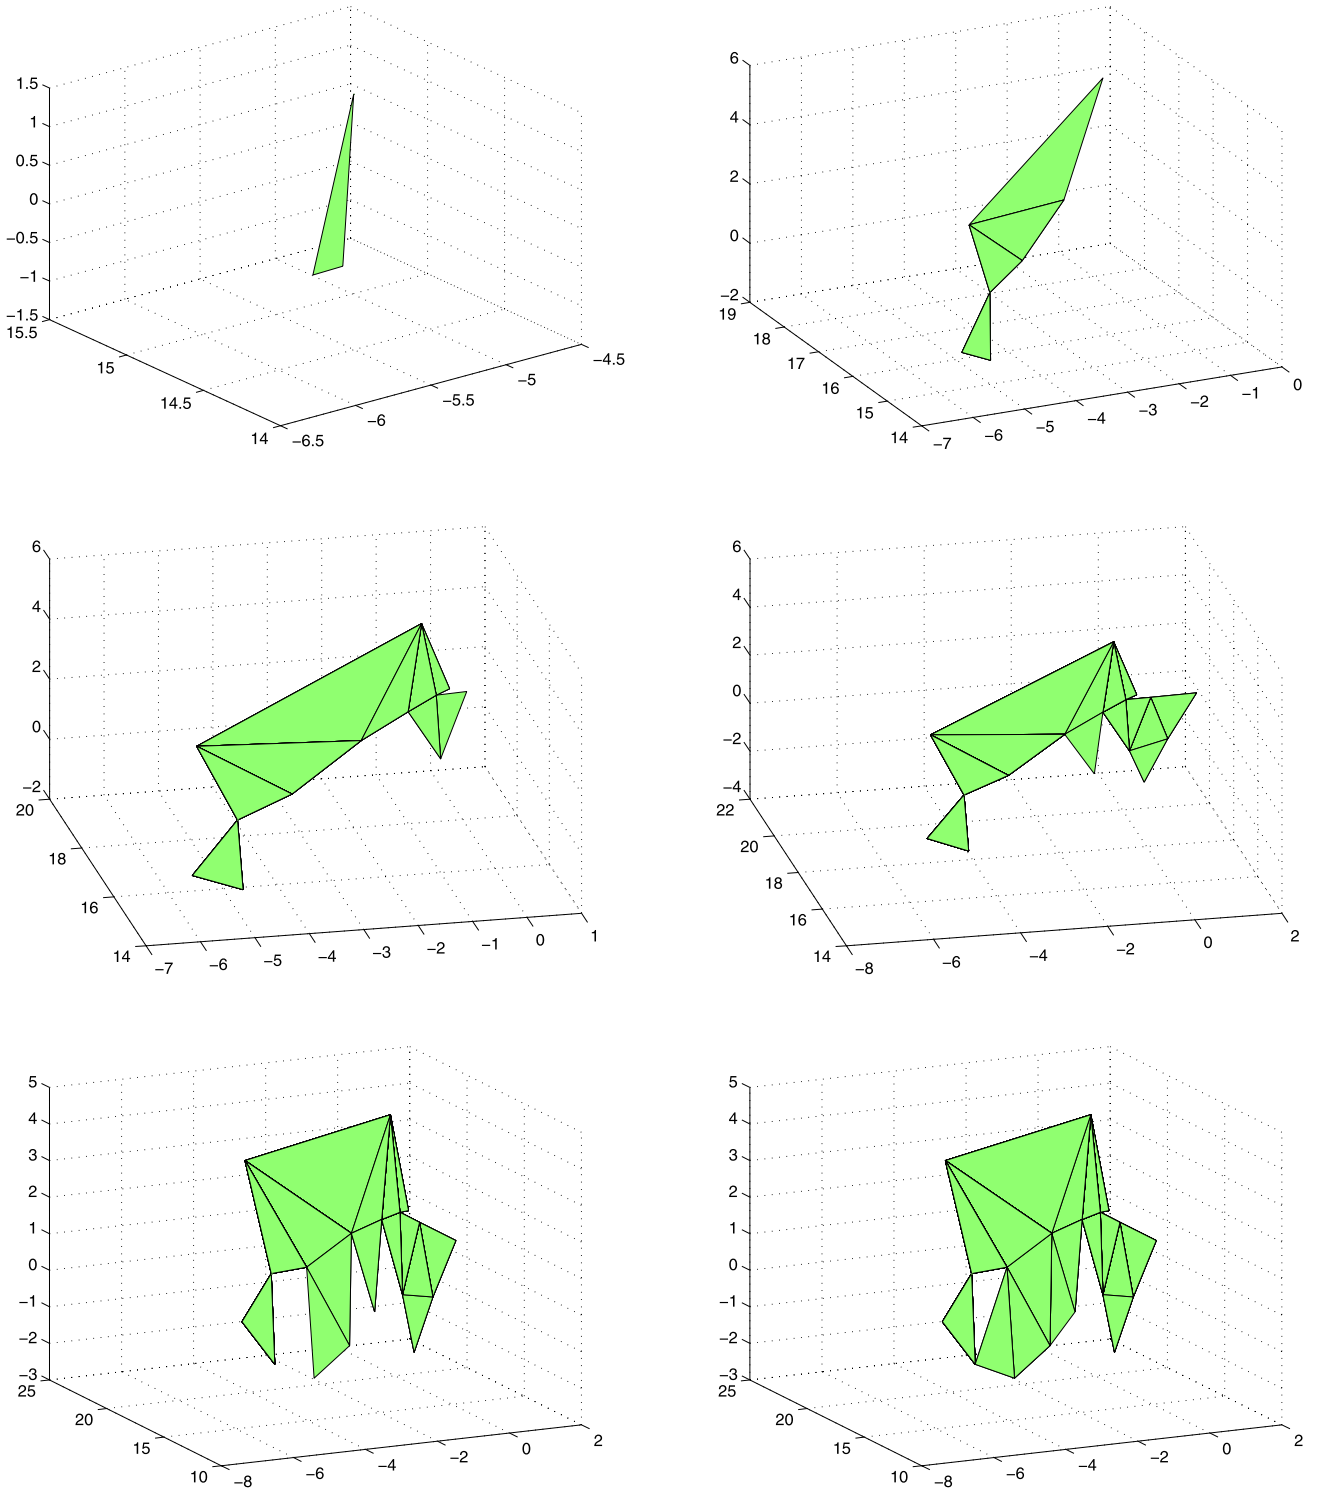

Fig. 10 A few stages in the building of a flat patch

All CT and MRI slices, both synthetic as well as real data, were farther processed through *3D Slicer* (www.slicer.org [25]), a software capable of various manipulations on 3D-scanned data, mainly for academic usage. We used *3D Slicer* in order to build triangulated brain/colon surfaces from sets of consecutive slices. Afterwards, both global/local algorithms presented herein were implemented on the triangulated surfaces using `Matlab`. Due to some limitations of

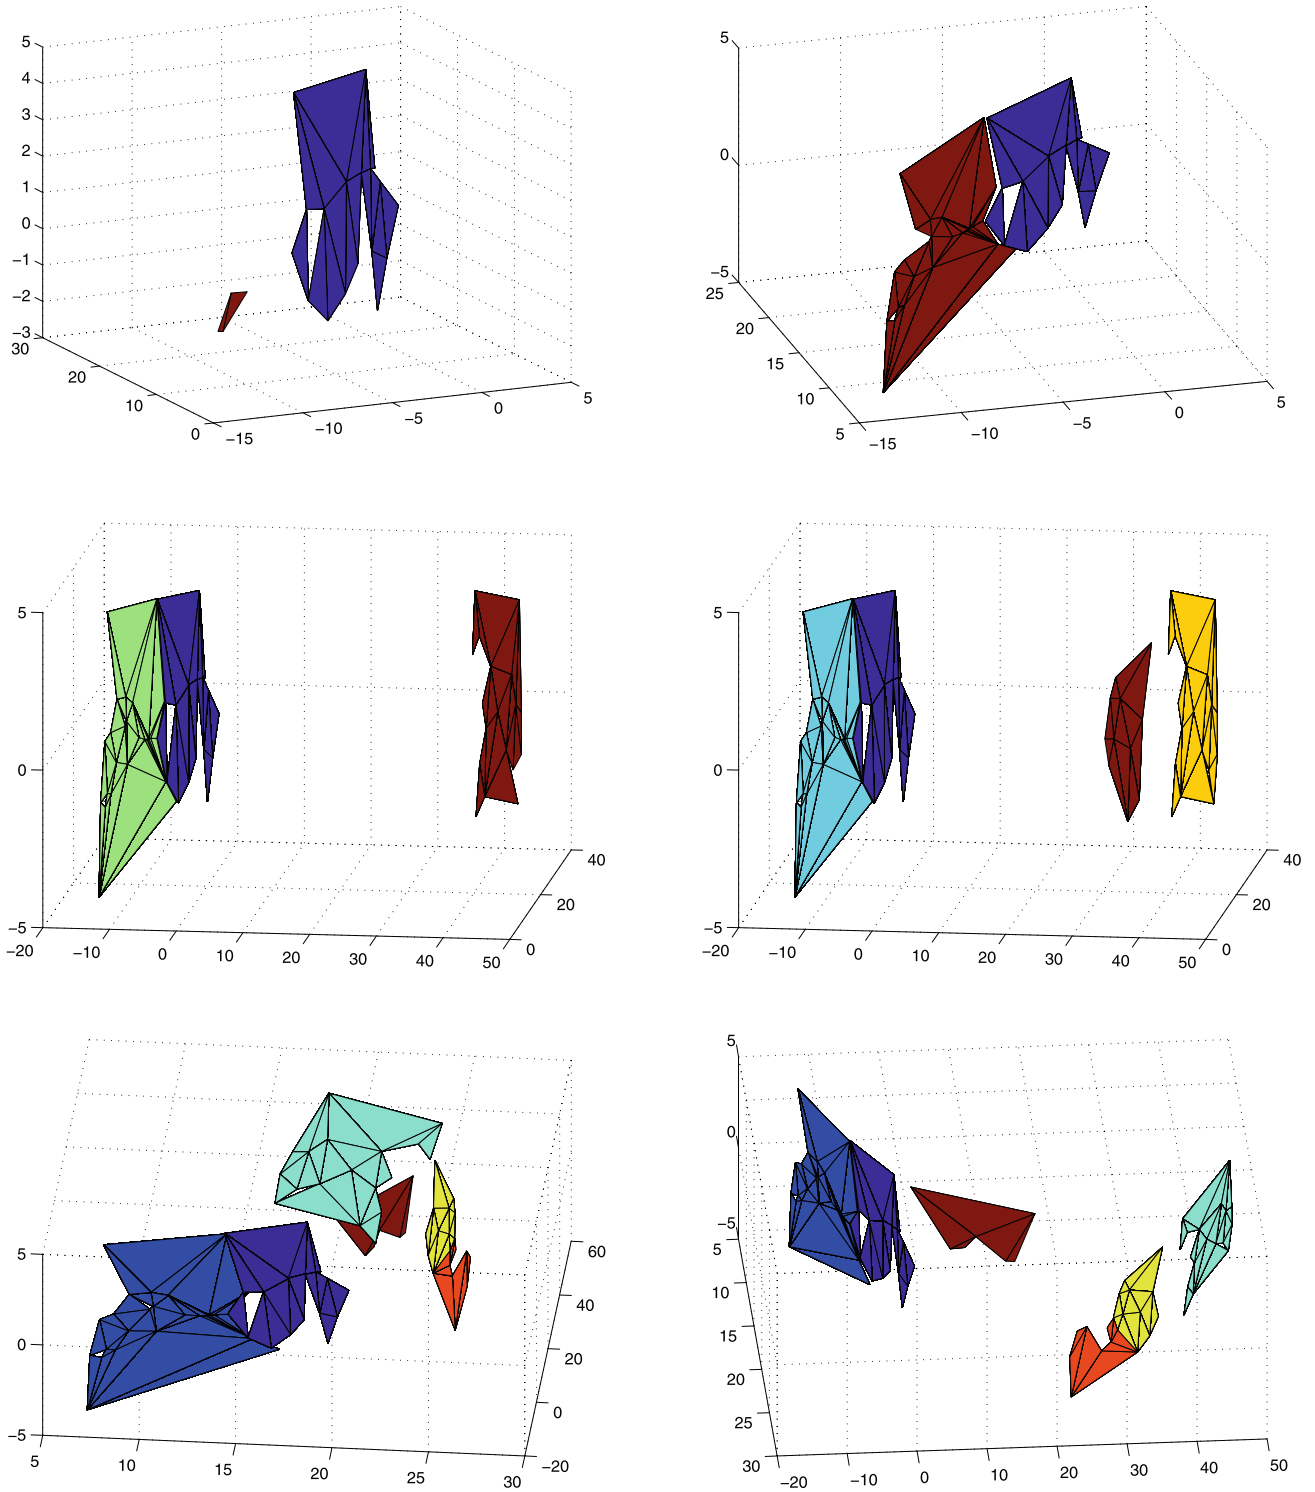

Fig. 11 Building the first patches of the flattening of half of the colon surface. Different view points (*bottom*) underlining the spacial distribution of the patches on an approximative half-cylinder

`Matlab` in processing and manipulating $3D$ graphic data, the colon slices were cut into two halves, denoted as the "front half" and, respectively, the "back half" of the colon piece on which algorithms were run.

4.1 Quasi-Geodesics and Conformal Modulus

This subsection focuses on the results obtained by implementing the algorithm for finding the straightest quasi-

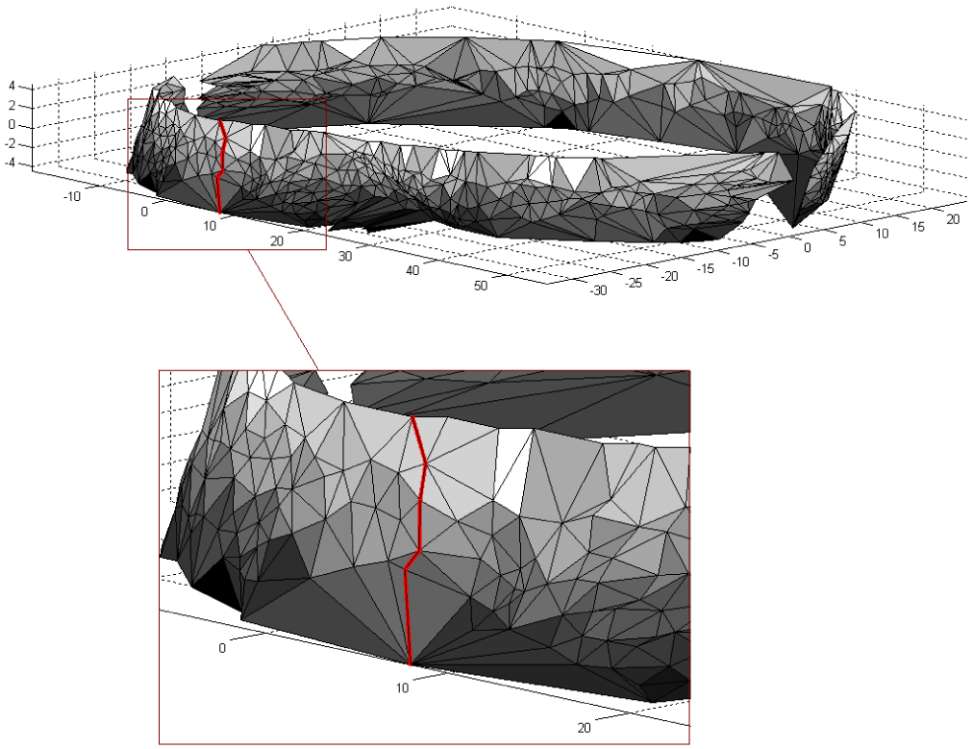

Fig. 12 (Color online) Shortest quasi-geodesic (*red line*) on the "back side" of the colon surface

geodesic path along a piece of the colon surface and the conformal modulus, computed with respect to this presumably best cut of the colon surface, in order to map it globally onto a flat region. The part of the colon shown below is resulted from 7 consecutive slices[2] that were processed through the `3D-Slicer` to give a triangulated surface.

4.1.1 Quasi-Geodesics

The straightest path from every possible vertex on the top end of the surface to the bottom end was found (see Fig. 12). The figure shows the shortest of all thus obtained quasi-geodesics. This "quasi-shortest" path is subsequently used in assessment of the conformal modulus. Its length is 9.82 whereas the Euclidean distance between its initial and final points as points in $\mathbb{R}^3$ is 9.38.

4.1.2 Conformal Modulus

Let γ be the shortest quasi-geodesic path found as above. Let its length be denoted by $L(\gamma)$. (As already mentioned above, $L(\gamma) = 9.82$.) After cutting the colon surface along γ, the resulting surface is a topological rectangle denoted by Q. Yet, its top/bottom sides may be of different lengths (see Fig. 10). Let $L_{\max}(L_{\min})$ denote the length of the longest (respectively shortest) of these. In the context of Sect. 2.2.3 the quasi-geodesic path γ corresponds to the path γ of (2.11) while for the circle c we employ $L_{\min}$.

Using (2.9) and (2.10) yields

$$\frac{L_{\min}^2}{Area(Q)} = \frac{175.68^2}{2051} = 15.04 \le M(Q) \le 21.26 = \frac{2051}{9.82^2} = \frac{Area(Q)}{L_\gamma^2}.$$

One can also resort to inequality (2.11) (or rather to its counterpart for quadrilaterals). For our purposes we shall compare the given irregular cylindrical surface C (i.e. colon slice) to the *mean cylinder* C_M, where the height and circumference of C_M are the mean height, 14, and circumference, 190, of C, respectively. By some elementary computations we find that the lower bound for $K = K_f$, in the case when $R \equiv C$ and $R' \equiv C_M$, is 1.5. While this appears to suggest that the mapping f is rather close to being conformal, recall that this is only a very crude estimate, based upon a rough averaging process. In fact, there are regions that can be mapped practically conformally, and others where the required dilatation is extremely high (see Sect. 4.2 below and Figs. 13, 14 and 15, therein).

4.2 Flattening Algorithm

We proceed to present and discuss experimental results obtained by applying the proposed algorithm in surface flattening, of both synthetic surfaces and data obtained from CT

[2] 7 slices being the largest manageable set, given `Matlab`'s memory limitations.

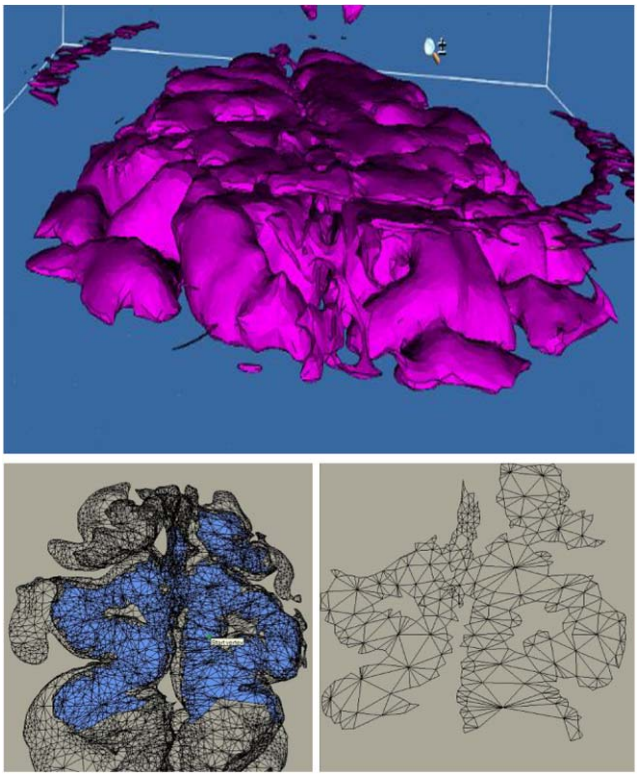

Fig. 13 Flattening of 3D Cerebral Cortex data obtained by MRI. The resolution is 15, 110 triangles, the chosen angle of 10° produces a dilatation of 1.1763. The location of the cortical region selected for flattening (*top*) and a non-simply connected patch (before flattening: *bottom*, *left* and after flattening: *bottom*, *right*). Notice the position of the starting vertex for the flattening algorithm

scans. Different GUIs, each presenting specific fine tunings, were employed. A fully functional improved GUI, including all the features of the previous versions and a number of additional ones, allowing the implementation of all the described algorithms is currently under construction and will be available soon.

In each of the examples both the input surface and a flattened representation of some patch are shown. Details about mesh resolution as well as flattening distortion can also be provided, as well as the number of patches needed in order to flatten the entire surface.

As already noted in Sect. 3, the algorithm was implemented in two processing modes—automatic and user defined.

The need for gluing patches together into a global image is clear (and is well known in Radiography as the "*pantomograph*").

However, here we do not address the problem of properly gluing patches. Therefore, in the example depicting the flattening and gluing patches, of half of the triangulated surface obtained from 3 CT-scan slices of the human colon, there are "holes" in the flattened representation, caused by artificially gluing neighboring patches to each other (see Fig. 15).

Evidently, as is apparent in the above examples of the colon flattening, one can have two neighbouring patches, with markedly different dilatations/distortions, which result in different lengths for the common boundary edges. Therefore, "cuts" and "holes" appear when applying a "naive" gluing. The discontinuities appear at the common bound-

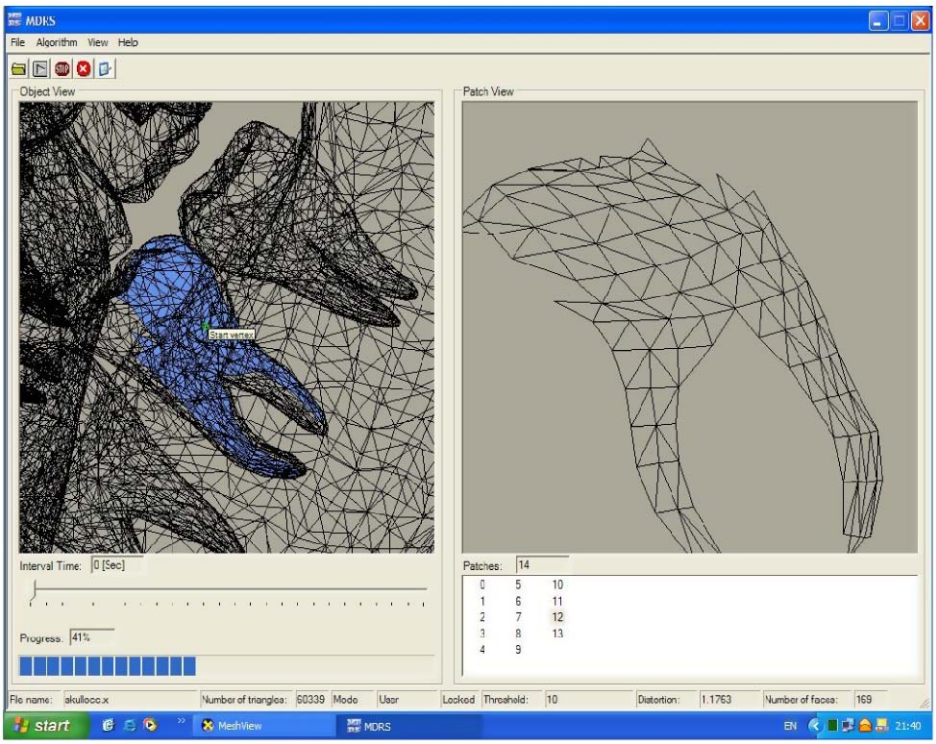

Fig. 14 Skull model—Detail: The resolution is of 60,339 triangles. $\alpha = 10°$ and the dilatation is 1.1763. Choosing the starting vertex in a region of low curvature results in large patches. Note that the GUI includes, amongst other features, the possibility to choose the angle α, hence the dilatation, which it computes and displays. It also shows the total number of patches and the number of triangles in each individual patch. Both user determined or random choice of the starting vertex are available

Table 1 Mesh resolution and patch size are given in number of triangles. "Patch size" refers to patches similar to the blue ones shown in Figs. 13 and 14

Object		Mesh resolution	Mode	#Patches	Patch size	Dilatation
Dolphin	(a)	564	auto	190	98	6.6713
	(b)	564	user	217	97	6.6713
Tooth	(a)	60339	auto	13	169	1.1763
	(b)	60339	auto	13	169	1.1763
Teapot	(a)	2256	auto	556	168	6.6713
	(b)	2256	auto	556	336	6.6713

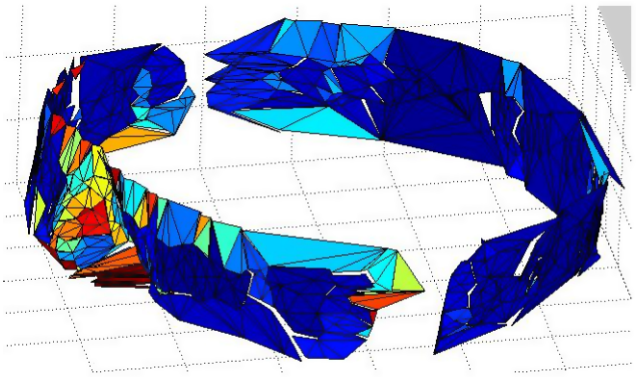

Fig. 15 Colon Flattening: Triangulated colon surface obtained from 3 CT-slices. Note the "cuts" and "holes" due to a "naive" gluing method. Each color represents different flattened patch. This and other results will be accessible through an interactive applet on the website [36]

ary of two patches obtained from regions with very different curvature. Indeed, the "back part" (Fig. 16 top) seems close enough to be half of a cylinder (and thus *developable*) but in fact it is highly folded and creased (see Fig. 16 bottom and Fig. 17).

In fact, the surface to be represented is topological cylinder of much higher curvature than the one in the visualization above, and with a "very bad" triangulation (Fig. 18, middle). This is a direct result of the scanning process, due to the difficulty to discern between colon and other tissues and between different folds of the colon (Fig. 18, bottom). This results in noisy data (Fig. 18, top) producing the "bad" triangulation observed.

Note that for small angles α (and hence for low dilatation/distortion) even non-simply connected patches may be obtained (e.g. Fig. 11 bottom, left). This is an immediate consequence of the algorithms' high sensitivity to curvature, quality that renders it highly suitable for discovering tumors, for instance in the colon, where they are associated with high curvature (Sect. 5).

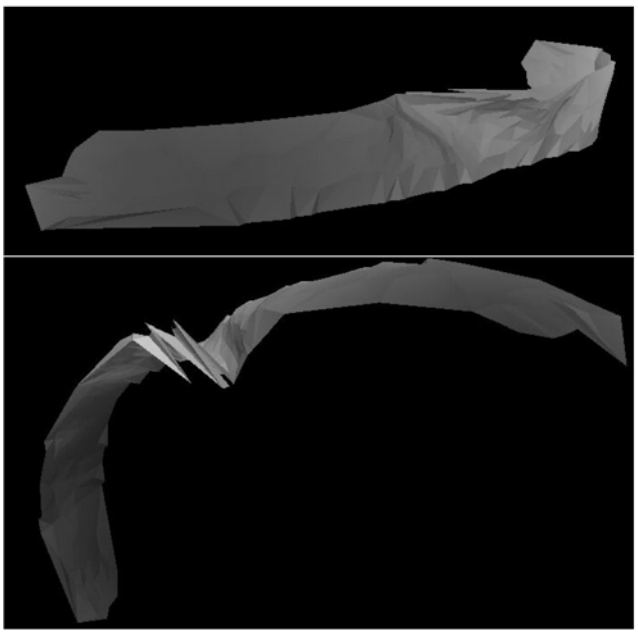

Fig. 16 Colon CT: Colon surface obtained from 7 slices of human colon scan. (Data curtesy of Dr. Doron Fisher from Rambam Medical Center in Haifa.) Frontal view (*top*) and rotated image (*bottom*) revealing the high degree of folding

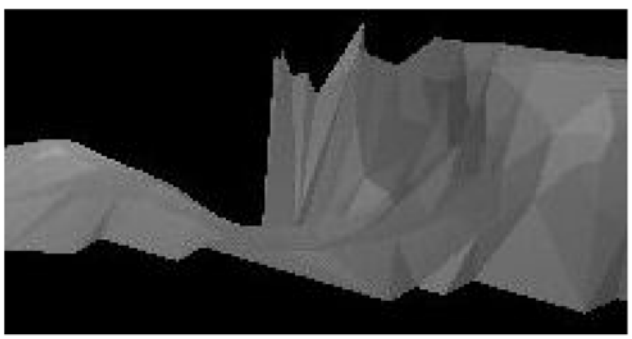

Fig. 17 Colon CT: Colon surface taken from 7 slices of human colon scan. Another detailed view of the fold in the colon surface

5 Concluding Remarks and Future Study

From the implementation results it is evident that this simple-to-program as well as efficient algorithm, also yields good flattening results and maintains small dilatations even in areas where curvature is large and good flattening is a challenging task. Moreover, since there is a simple way to assess the resulting dilatation, the algo-

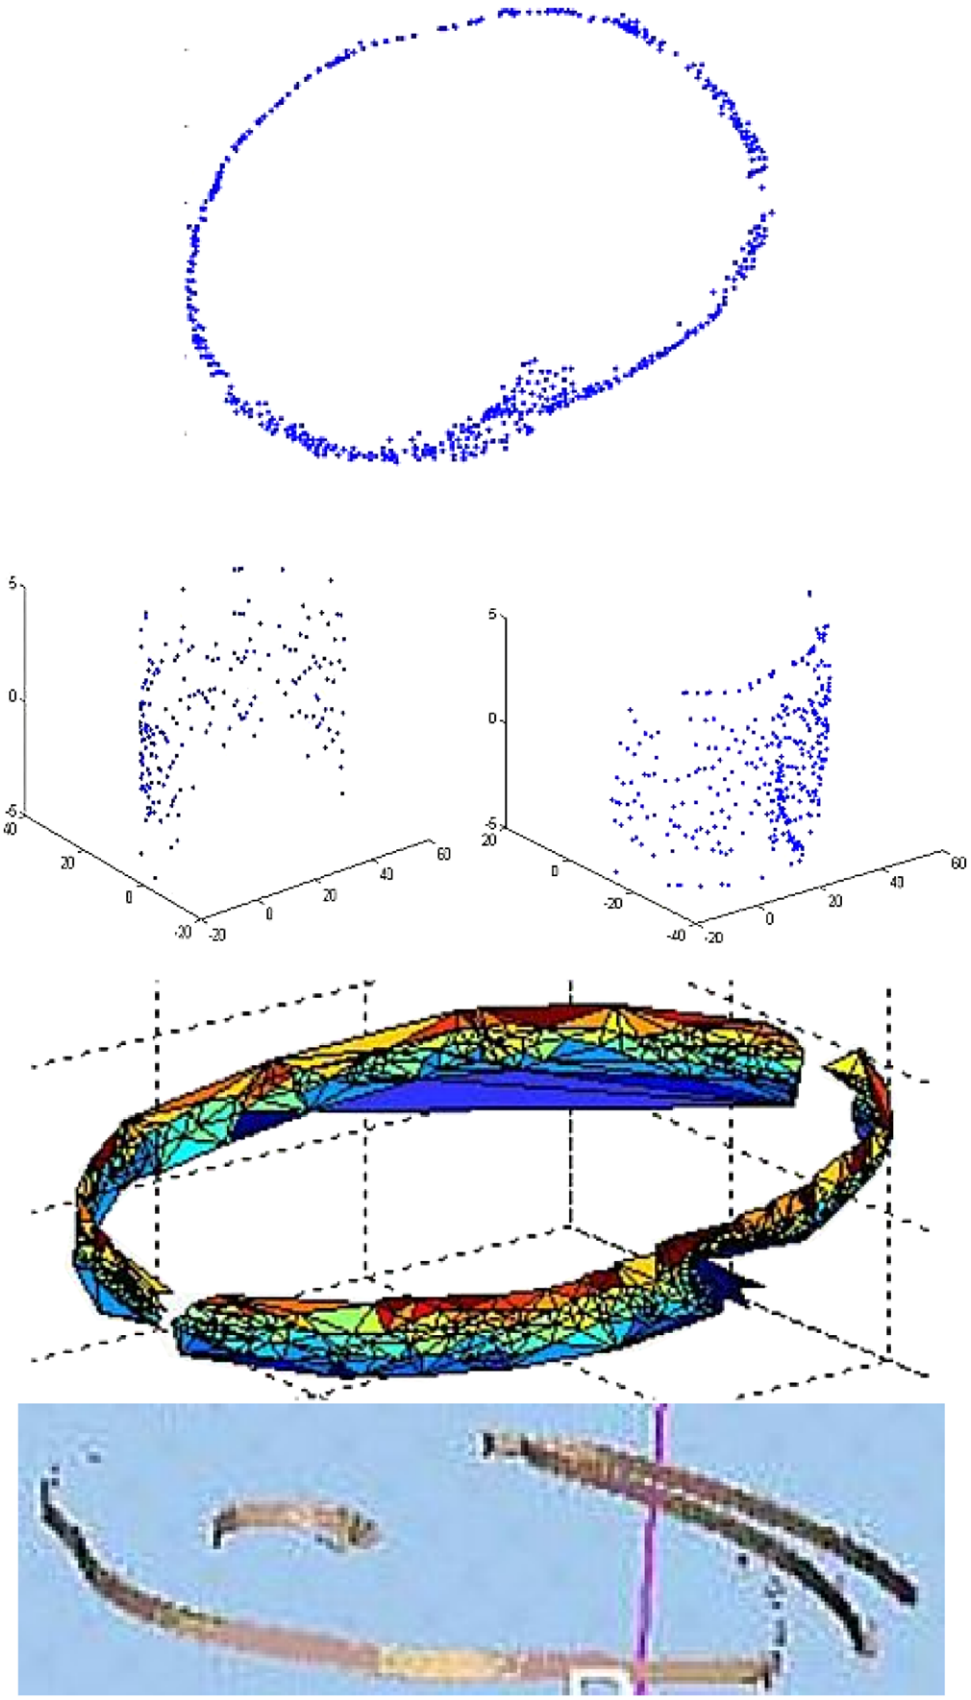

Fig. 18 Colon CT: Noisy sampling (*top*) producing an uneven triangulation (*middle*) due to the difficulty in discerning the targeted tissue (*bottom*)

rithm was implemented in such a way that the user can set in advance an upper bound on the resulting dilatation.

Contrary to some related studies, our algorithm does not require the use of derivatives. Consequently, the algorithm does not suffer from typical drawbacks of derivative-based computations. Moreover, since no derivatives are employed, no smoothness assumption about the surface to be flattened are made, which makes the algorithm presented herein ideal for use in cases where smoothness is questionable.

However, for the existence of faithful quasi-conformal/quasi-isometric representations for sampled surface (and higher dimensional manifolds) strongly depends on the quality of the sampling and reconstruction, i.e. basically on the possibility of finding good *PL* approximations to smooth manifolds. (See [29] for a detailed study along these lines.) The largest distortion/dilatation will be yielded by *cone points* of high discrete curvature (see [21] for a discussion, albeit based upon a different approach, on this type of singularity). In the frame presented herein, one can cope with this problem in either one of the following ways: (a) Consider the surfaces as obtained, without any preprocessing, and view the ensuing small patches as an indicator of bad curvature and folding (a consequence, perhaps, of "bad" sampling); or (b) Adopt the approach familiar in Signal and Image Processing and smoothen the surface to obtain, in a controlled manner, regions of concentrated lower curvature. (We thank Prof. Peter Maas for suggesting us this method.)

The algorithm may be practical for application in cases where local yet good analysis is required, in medical imaging, emphasizing *accurate* and *measurable* flattened representation of the brain and the colon (virtual colonoscopy). We have noted in the previous section the algorithm's high sensitivity to curvature and the advantage this provides in the detection of carcinogenic polyps in the colon. In this context, it is important to underline here the fact that quasi-conformal mappings, via the dilatation, are highly sensitive to even *local* variations in the geometry of the surface under analysis. This renders them ideal to discern in such minute variations, as produced by polyps and that may be smoothed out (by flow methods) or discarded as noise (by statistical methods). In addition to its CAD (Computer Aided Detection) capability, another possible application of our algorithm would be in the mapping of the human brain, since high curvature regions are usually associated with more developed regions, and hence, for instance, to higher intelligence. (We are thankful to Prof. Jacob Rubinstein who suggested this possible application.) This type of application appears to be useful in Comparative Anatomy and Paleontology. Such directions of study are currently explored.

Further research is required however in processing from local to flat global image in a more precise fashion, i.e. how can one glue two neighbouring patches while keeping fixed bounded dilatation (i.e. to actually compute the *holonomy* of the surface—see [34]). Indeed, we may flatten the neighborhood of two vertices u, v, obtaining the flat images I_u, I_v, so that these two neighborhood have some intersection along the boundary yet, it will not be possible to adjust the resulting images to give one flat image $I_u \cup I_v$ of the union of these neighborhood which still keeping the quasi-isometric property. Here too, study is underway.

It should also be noted that, while the application presented here is for $2D$-images of $3D$-surfaces, the results of Gehring and Väisälä are stated and proven for any dimension. Therefore, implementations for higher dimensions are also underway (see, for instance, [28] for some possible research directions).

Indeed, since the importance of studying the elastic properties of images (i.e. quasi-isometries), as opposed to the rigid ones (i.e. isometries) has been already stressed in [23] (see also [35]), and since planar quasi-isometries almost preserve shape when strain is under the critical value [9], the application of quasi-isometries in higher-dimensional image processing and related fields is necessary, particularly so in Computerized Tomography.

An additional promising direction of study stems from the theory of quasi-isometric/quasi-conformal mappings between admissible surfaces, with applications in face recognition.

Acknowledgements Emil Saucan was partly supported by the Viterbi Postdoctoral Fellowship. Research partly supported by the Ollendorf Minerva Center.

The authors would like to thank the reviewers for their patience and their many helpful comments and suggestions.

Thanks are also due to Assaf Manor, Pavel Mikhlin and Ofir Zeitoun for their skillful programming that helped produce some of the images included herein.

References

1. Alexandrov, A.D., Reshetnyak, Y.G.: General Theory of Iregular Curves. Mathematics and its Applications, vol. 29. Kluwer Academic, Dordrecht (1989)
2. Angenent, S., Haker, S., Tannenbaum, A., Kikinis, R.: On area preserving maps of minimal distortion. In: Djaferis, T., Schick, I. (eds.) System Theory: Modeling, Analysis, and Control, pp. 275–287. Kluwer Academic, Dordrecht (1999)
3. Appleboim, E., Saucan, E., Zeevi, Y.: Quasi-Conformal Flat Representation of Triangulated Surfaces for Computerized Tomography, CVAMIA 2006. Lecture Notes in Computer Science, vol. 4241, pp. 155–165. Springer, Berlin (2006)
4. Berger, M.: A Panoramic View of Riemannian Geometry. Springer, Berlin (2003)

5. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: On isometric embedding of facial surfaces into S3. In: Proc. Intl. Conf. on Scale Space and PDE Methods in Computer Vision, pp. 622–631 (2005)
6. Caraman, P.: n-Dimensional Quasiconformal (QCf) Mappings. Abacus Press, Tunbridge Wells (1974). Editura Academiei Române, Bucharest
7. do Carmo, M.P.: Differential Geometry of Curves and Surfaces. Prentice-Hall, Englewood Cliffs (1976)
8. Clarenz, U., Litke, N., Rumpf, M.: Axioms and variational problems in surface parameterization. Comput. Aided Geom. Des. **21**(8), 727–749 (2004)
9. Gehring, W.F.: Topics in quasiconformal mappings. In: Vuorinen, M. (ed.) Quasiconformal Space Mappings—A Colection of Surveys 1960–1990. Lecture Notes in Mathematics, vol. 1508, pp. 20–38. Springer, Berlin (1991)
10. Gehring, W.F., Väisälä, J.: The coefficients of quasiconformality. Acta Math. **114**, 1–70 (1965)
11. Gu, X., Wang, Y., Yau, S.T.: Computing conformal invariants: period matrices. Commun. Inf. Syst. **2**(2), 121–146 (2003)
12. Gu, X., Yau, S.T.: Computing conformal structure of surfaces. Commun. Inf. Syst. **2**(2), 121–146 (2002)
13. Gu, X., Yau, S.T.: Global conformal surface parameterization. In: Eurographics Symposium on Geometry Processing (2003)
14. Haker, S., Angenet, S., Tannenbaum, A., Kikinis, R.: Non distorting flattening maps and the 3-D visualization of colon CT images. In: IEEE Transauctions on Medical Imaging, vol. 19, No. 7 (2000)
15. Haker, S., Angenet, S., Tannenbaum, A., Kikinis, R., Sapiro, G., Halle, M.: Conformal surface parametrization for texture mapping. In: IEEE Transauctions on Visualization and Computer Graphics, vol. 6, No. 2 (2000)
16. Hallinan, P.: A low-dimensional representation of human faces for arbitrary lighting conditions. In: Proc. CVPR, pp. 995–999 (1994)
17. Hilbert, D., Cohn-Vossen, S.: Geometry and the Imagination. Chelsea, New York (1952)
18. Hurdal, M.K.: Personal communication to ES
19. Hurdal, M.K., Bowers, P.L., Stephenson, K., Sumners, D.W.L., Rehm, K., Schaper, K., Rottenberg, D.A.: Quasi conformally flat mapping the human crebellum. In: Taylor, C., Colchester, A. (eds.) Medical Image Computing and Computer-Assisted Intervention—MICCAI'99, vol. 1679, pp. 279–286. Springer, Berlin (1999)
20. Hurdal, M.K., Stephenson, K.: Cortical cartography using the discrete conformal approach of circle packings. NeuroImage **23**, 119–128 (2004)
21. Kharevich, L., Springborn, B., Schröder, P.: Discrete conformal mapping via circle patterns. ACM Trans. Graph. **25**(2), 412–438 (2006)
22. Kimmel, R., Malladi, R., Sochen, N.: Images as embedded maps and minimal surfaces: movies, color, texture, and volumetric medical images. Int. J. Comput. Vis. **39**(2), 111–129 (2000)
23. Koenderink, J.J.: Solid Shape. MIT Press, Cambridge (1990)
24. Lehto, O., Virtanen, K.I.: Quasiconformal Mappings in the Plane. Springer, Berlin (1973)
25. Pieper, S., Halle, M., Kikinis, R.: 3D Slicer. In: IEEE International Symposium on Biomedical Imaging: Nano to Macro, vol. 1, pp. 632–635. IEEE Society (2004)
26. Polthier, K., Schmies, M.: Straight geodesics on polyhedral surfaces. In: Hege, H.C., Polthier, K. (eds.) Mathematical Visualization, pp. 135–150. Springer, Berlin (1998)
27. Pogorelov, A.V.: Quasigeodesic lines on a convex surface. Am. Math. Soc. Transl. **6**(72), 430–472 (1952)
28. Saucan, E., Appleboim, E., Zeevi, Y.Y.: Image projection and representation on S^n. J. Fourier Anal. Appl. **13**(6), 711–727 (2007)
29. Saucan, E., Appleboim, E., Zeevi, Y.Y.: Sampling and reconstruction of surfaces and higher dimensional manifolds. J. Math. Imaging Vis. **30**(1), 105–123 (2008)
30. Seung, H.S., Lee, D.D.: The manifold ways of perception. Science **290**, 2323–2326 (2000)
31. Sheffer, A., de Stuler, E.: Parametrization of faceted surfaces for meshing using angle based flattening. Eng. Comput. **17**, 326–337 (2001)
32. Stephenson, K.: Personal communication to ES
33. Surazhsky, T., Magid, E., Soldea, O., Elber, G., Rivlin, E.: A comparison of Gaussian and mean curvatures estimation methods on triangular meshes. In: Proceedings of the IEEE International Conference on Robotics and Automation, Taipei, Taiwan, pp. 1021–1026 (2003)
34. Thurston, W.: Three-Dimensional Geometry and Topology, vol. 1. Princeton University Press, Princeton (1997). (S. Levy (ed.))
35. Wardetzky, M., Bergou, M., Harmon, D., Zorin, D., Grinspun, E.: Discrete quadratic curvature energies. Comput. Aided Geom. Des. **24**(8–9), 499–518 (2007)
36. http://visl.technion.ac.il/ImageSampling
37. Zayer, R., Lévy, B., Seide, H.-P.: Linear angle based parameterization. In: Belyaev, A., Garland, M. (eds.) Proceedings of Eurographics Symposium on Image Processing (2007)

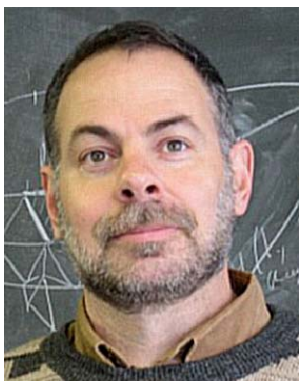

Emil Saucan is a Senior Research Fellow at the Electrical Engineering and Mathematics Department, Technion–Israel Institute of Technology. He has received his BSc in Pure Mathematics from the University of Bucharest and his MSc and PhD, also in Pure Mathematics from the Technion. He subsequently was a Viterbi postdoctoral fellow at the Electrical Engineering Department, Technion. While his Ph.D. Thesis formally belongs to the fields of Geometric Function Theory and Discrete Groups, his main research interest is Geometry in general (including Geometric Analysis and Geometric Topology), especially Discrete and Metric Differential Geometry (and its applications to Imaging and Geometric Design) and Geometric Modeling. Currently his efforts are concentrated in the application of these fields, to Computer Graphics, Medical Imaging, Signal and Image Sampling and Reconstruction, Bio-Informatics and Communication Networks. He has also published papers on Mathematics Education, both Pure and Applied, mainly on diverse methods and approaches of teaching Geometry, Topology and their applications.

Eli Appleboim Eli Appleboim is a lab engineer and research assistant, at The Vision and Image Sciences Laboratory, Faculty of Electrical Engineering, Technion–Israel Institute of Technology. He has received his BSc and MSc in Pure Mathematics, both from the Technion. He has worked as a senior software/algorithm engineer at Broadcom, Israel (VisionTech), being involved in the design of algorithms for MPEG2 video processing. His main research interests are Computer vision, Image processing, Computerized Tomography and Low Dimensional Topology and Geometry, mainly applications of 3-Manifolds and Knot theories to Image Processing.

Efrat Barak-Shimron is a MSc student in the Interdisciplinary Program in Neuroscience, Faculty of Medicine, Technion–Israel Institute of Technology. She received a dual BSc in Physics and a BA in Electrical Engineering, both *cum laude*, also from the Technion. She has been awarded a number of prizes awards, amongst them the Irwin and Joan Jacobs Graduate School scholarship and the Presidents and Dean's Excellence Awards, Technion. Her current subject of scientific research is Topology and Synchronization in Neural Networks.

Ronen Lev is a Senior Software Engineer at DAZ Productions, Draper, Utah. He has received his BA in Software Engineering from Technion–Israel Institute of Technology and his MSc in Computer Science from the University of Haifa, Israel. He has worked as a Software team leader at OptiTex, Israel. Before that he was an Advanced Algorithm Officer at IDF and a Software developer at Rafael, Israel. His main fields of interests are Computer Graphics, the development and programming of algorithms for applications of Differential Geometry, CAD/CAM Applications and Theory-meets-Practice solutions.

Yehoshua Y. Zeevi is the Barbara and Norman Seiden Professor of Computer Sciences in the department of Electrical Engineering, Technion–Israel Institute of Technology. He is the Director of the Ollendorff Center for Vision and Image Sciences, and the Head of the Zisapel Center for Nano-Electronics. He received his Ph.D. from the University of California, Berkeley, and was subsequently a Vinton Hayes Fellow at Harvard University, where he had been a regular visitor for many years. He was also a Visiting Professor at MIT, the CAIP Center of Rutgers university, and Columbia University. His major research has been devoted to vision and image sciences, and to signal and image representation and processing. He is a Fellow of the SPIE and the Rodin Academy, the Editor-in-Chief of the Journal of Visual Communication and Image Representation, published by Elsevier. He is one of the founders of i Sight, Inc.,—a company devoted to digital photography and real time image processing, and of UltraGuide—a company that developed innovative computerized systems for visual guidance in ultrasound-assisted medical procedures, and of Cortica—a company that developed new large scale systems for tracking and deep content classification. He is a member of the IEEE Committee for Multidimensional Signal processing and of the Board of the Technion.