Computers & Graphics 37 (2013) 582–588

Contents lists available at SciVerse ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag

SMI 2013

Generalized extrinsic distortion and applications

Patricio Simari [a,b,*], Leila De Floriani [a,b], Federico Iuricich [a,b], Mohammed M. Mesmoudi [a,b]

[a] Department of Computer Science, University of Maryland, A.V. Williams Bldg, College Park, MD 20740, United States
[b] DIBRIS, University of Genova, Via Dodecaneso n 35, 16146 Genova, Italy

ARTICLE INFO

Article history:
Received 18 March 2013
Received in revised form
29 May 2013
Accepted 29 May 2013
Available online 14 June 2013

Keywords:
Geometry processing
Discrete mean curvature
Distortion

ABSTRACT

Curvature is a key feature in shape analysis and its estimation on discrete simplicial complexes benefits many geometry processing applications. However, its study has mostly remained focused on 2D manifolds and computationally practical extensions to higher dimensions remain an active area of computer science research. We examine the existing notions of distortion, an analog of curvature in the discrete setting, and classify them into two categories: *intrinsic* and *extrinsic*, depending on whether they use the interior or the dihedral angles of the tessellation. We then propose a generalization of extrinsic distortion to ce:italic> D /ce:italic> D and derive a weighting that can be used to compute mean curvature on tessellated hypersurfaces. We analyze the behavior of the operator on 3-manifolds in 4D and compare it to the well-known Laplace–Beltrami operator using ground truth hypersurfaces defined by functions of three variables, and a segmentation application, showing it to behave as well or better while being intuitively simple and easy to implement.

1. Introduction

The local curvature of a surface is very descriptive and an important tool in geometry processing. The applications of its estimation to the analysis of discrete surfaces have been extensively studied, including mesh simplification [1,2], alignment [3], ridge-valley line detection [4], non-photorealistic rendering [5], segmentation [6,7], partial shape matching [8], symmetry detection [9], denoising [10,11], and remeshing [12].

Curvature estimation methods can be broadly classified into two categories: fitting methods and discrete methods. The former use local regression to estimate the parameters of continuous models and evaluate curvature using its continuous definition. The latter find discrete analogues to the continuous elements involved in defining curvature so that the notion can be evaluated directly in the discrete domain.

This work will focus on discrete methods. While fitting methods are more tolerant to noise and tessellation artifacts, they can be more computationally intensive. This is a drawback in applications that require curvature to be estimated in real time, *e.g.*, physical simulation, non-photorealistic rendering, or real-time shape analysis for robotics applications. Discrete methods are simple to implement, require fewer computations, and are trivially parallelizable.

A volume dataset can be seen as a 3-manifold hypersurface embedded in 4D and, as such, is amenable to 3D curvature analysis. Similarly, such analysis can be conceived for hypersurfaces in 4D space which are not the graph of a 3D scalar field, such as isosurfaces of time-varying scalar fields or tetrahedral meshes defined by animation sequences [13]. However, the application of curvature and its discrete variants to volumetric shape analysis remains comparatively unexplored. One exception is the concept of discrete distortion.

We will discuss the existing definitions of distortion for discrete surfaces, classify them into two categories, *intrinsic* and *extrinsic*, and present a generalization of extrinsic distortion to n D, deriving a weighting that can be used to compute mean curvature. We will analyze the behavior of the operator on 3-manifolds in 4D and compare it to the well known Laplace–Beltrami operator in two ways. First, we examine the behavior on a suite of analytic surfaces sampled under varying conditions of resolution, distribution of samples, and noise. Second, we examine using the distortion field to obtain volumetric segmentations and evaluate their stability under increasing image noise. We will show, in each case, that extrinsic distortion behaves similarly or better than the Laplace–Beltrami operator while being intuitively simple and easy to implement.

2. Related work

We will assume familiarity with the fundamental notions of curvature in the continuum and refer the reader to the relevant

* Corresponding author.
E-mail addresses: pdsimari@gmail.com, psimari@cs.umd.edu(P. Simari)

0097-8493/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.cag.2013.05.018

background [14]. In the context of discrete representations of 2D surfaces such as meshes and point clouds, curvature is a well-studied area due to its numerous applications [15,16]. There are a plethora of methods for curvature evaluation, each with advantages and disadvantages. These methods can be broadly classified into *fitting* methods and *discrete* methods.

Fitting methods use local regression to fit a continuous function, such as a polynomial, to the surface data near a point of interest. Once the parameters of such a function are determined, a curvature estimate can be computed analytically or using finite element methods [15–20]. Other approaches, rather than fitting a model to surface points, fit the curvature tensor to normal variations in a local neighborhood [5,11]. In contrast, discrete approaches [18,21,10,22–24] compute quantities that approximate curvature values directly on the discrete surface without explicitly fitting a continuous model.

Within the context of 3-manifolds in 4D, there is much less work within the geometry processing community. Hamann introduces a generalization of the polynomial fitting approach to extend to such data [25], but this is based on a continuous method rather than a discrete one. Within the discrete setting, the notion of discrete distortion arises as a purely discrete analog to continuous curvature [26]. It has been successfully used in several applications, including morphological analysis [27], guiding multi-resolution simplification [28], and medical visualization [29].

In the above cases, this particular notion of discrete distortion is an intrinsic measure based on a generalization of concentrated curvature [30,31] and angle deficit [10]. In 2D, however, the previously introduced notion of distortion is based on dihedral angles and is related to mean curvature [32] and, as such, is an extrinsic measure. While the distinction between intrinsic and extrinsic measures is a well known one in the context of curvature, this paper applies these notions to categorize the different previously published notions of distortion and unify them. Moreover, we introduce here a generalization of extrinsic distortion to n D.

3. Background

3.1. Intrinsic distortion in 3D

For tetrahedral meshes embedded in four-dimensional Euclidean space a separate notion of distortion has been introduced [26], in this case, as a generalization of Aleksandrov's concentrated curvature [30] to higher dimensions, which can be considered a discrete counterpart to the scalar Ricci curvature [33]. A similar approach has been independently proposed [34]. This is an *intrinsic* measure based on angle deficits, which in the 2D case constitutes a discrete counterpart to Gaussian curvature. In the 3D case, given a tetrahedralized manifold embedded in 4D space, the intrinsic distortion at an internal vertex p can be defined as

$$D(p) = 4\pi - \sum_{i=1}^{n} S_i,$$

where S_i is the solid angle at of the i-th tetrahedron incident on p at said vertex.

3.2. Extrinsic distortion in 2D

In the 2D case, distortion has been previously defined at a vertex [32] as follows. The idea is to compute the sum of the angle deficits of the dihedral angles at the edges incident on p, with respect to the flat angle. Vertex distortion at an internal vertex p of the triangulation is defined as

$$D(p) = \sum_{i=1}^{N} (\pi - \Theta_i),$$

where $e_1, \ldots, e_N$ are the edges incident on p, and Θ_i is the dihedral angle formed between the two triangles incident at edge e_i.

A weighted version of this formulation can be used to estimate mean curvature. This weighted form coincides with the cylindrical approximation method [35]. For each edge, the integral form is obtained by weighting the angle by half of the edge length for each vertex in the edge. Another factor of one half is introduced by the fact that the one of the principal curvatures of a cylinder is null, thus causing the mean curvature to be one half that of the non-null principal curvature. The final punctual form of the mean curvature estimate is obtained by dividing the area A_p associated with the vertex p, computed as the sum of fractional areas of all triangular faces incident on p. This fraction can be taken to be a fixed 1/3, leading to the barycentric formulation, or the Voronoi region can be used instead. The final expression of the punctual form is thus

$$\hat{H}(p) = \frac{1}{4A_p} \sum_{i=1}^{N} \|e_i\| (\pi - \Theta_i).$$

3.3. The Laplace–Beltrami operator

One of the best known tools for the computation of mean curvature on discrete surfaces is the Laplace–Beltrami operator [10]. It can be readily generalized to 3D manifolds and, as a well-known discrete operator and estimator of mean curvature, we will use it as a basis for comparison of our operator in the experimental section. Its value at a vertex p is given by

$$\mathbf{K}(p) = \frac{1}{V(p)} \sum_{i \in N_1(p)} w_i (p - x_i)$$

In 3-manifolds, the mean curvature value is given by $\frac{1}{3}\|\mathbf{K}\|$. Here, $V(p)$ denotes the tetrahedral volume assigned to vertex p and w_i denotes the weight associated with the edge (p, x_j). In 3D, this weight is given by

$$w_i = \frac{1}{6} \sum_j \ell_i^j \cot \alpha_i^j$$

where ℓ_i^j is the length of the edge *opposite* to edge (p, x_i) within the j-th tetrahedron incident on (p, x_i), and α_i^j is the dihedral angle at this opposite edge.

For the value of the volume $V(p)$ we simply use barycentric volumes, obtained as 1/4th of the sum volume of all tetrahedra incident on p. Alternatively, it is possible to use Voronoi volumes, though we have found that, in the 3D case, it does not reliably improve the accuracy of the mean curvature estimate.

The scalar value of the operator as defined above would always be a positive quantity, given that it is a fraction of the norm of the $\mathbf{K}$ vector. However, we can set the sign of the scalar value by setting it to match the sign of the dot product between $\mathbf{K}$ and the manifold normal (positive when they are in agreement, negative when opposite). In the case of graphs of scalar fields, as we will be examining, we can simply use the sign of the last component of $\mathbf{K}$ for efficiency and simplicity.

We consider the Laplace–Beltrami operator because of how well established it is in the literature as a discrete estimator of mean curvature. A thorough comparison of a large set of curvature operators is beyond the scope of this work and, for such a comparison, we refer the reader to published works [11,36]. Our main goal is to compare our approach to a well known one, establishing it to behave as well as said approach with some advantages, and thus placing it in context.

4. Generalizing extrinsic distortion to *n* D

Here, we wish to generalize this *extrinsic* notion of distortion from 2D to higher dimensions. We can naturally do so by considering the dihedral angle at adjacent simplexes. On a discrete n-dimensional manifold, embedded in $(n+1)$D, represented by a simplicial complex Σ, pairs of adjacent n-simplexes form a dihedral angle determined by the two hyperplanes containing each of them. Assuming that the manifold is orientable, the signed dihedral angle formed by these hyperplanes can be determined in a straightforward manner, leading us to formulate the general expression for extrinsic distortion:

$$D(p) = \sum_{\tau_{ij} \in \mathrm{St}^2(p)} (\pi - \Theta_{ij}),$$

where Θ_{ij} represents the signed dihedral angle between the simplexes σ_i and σ_j, and $\tau_{ij} \in \mathrm{St}^2(p)$ is defined as true if $\sigma_i, \sigma_j \in \mathrm{St}(p)$ and $\tau_{ij} = \sigma_i \cap \sigma_j$, where τ_{ij} is an $(n-1)$-simplex. This is to say σ_i and σ_j are adjacent and their union has disk topology.

4.1. Mean curvature in n D

The weighting that leads to a mean curvature approximation can also be generalized, inferring it from Dyn's cylindrical approximation in 2D [35]. Two adjacent n-simplexes will meet at an $(n-1)$-simplex τ_{ij}, and the edge length used in the 2D case can be generalized to the volume of τ_{ij}. Just as in the 2D case one half of the weighted angle went to each vertex on the edge, in the general case $1/n$ goes to each of the vertices at the simplicial intersection. Finally, the cylindrical generalization has $n-1$ null principal curvatures and thus its mean curvature is given by $1/n$-th of the non-null value. These all lead to the final weighted expression:

$$\hat{H}(p) = \frac{1}{n^2 \|p\|} \sum_{\tau_{ij} \in \mathrm{St}^2(p)} (\pi - \Theta_{ij}) \|\tau_{ij}\|,$$

where $\|p\|$ is the n-dimensional barycentric volume associated with vertex p and $\|\tau_{ij}\|$ is the $(n-1)$-dimensional volume associated with the simplex τ_{ij} at the intersection of the adjacent simplexes σ_i and σ_j. In the particular case of a 3-manifold, $n=3$, the simplexes are tetrahedra, adjacent simplexes meet at triangles, and dihedral angles are thus weighted by triangle area.

Derivation: In the following we provide a derivation of the above weighting, which is not intended as a proof of convergence. First, let us remark that when two hyperplanes in R^{n+1} intersect, the intersection is an $(n-1)$-affine plane $\mathcal{P}_{n-1}$. We can approximate smoothly the PL-hypersurface through a cap of the curved nD hypersurface $\mathcal{C}_n(r) = S^1_r \times \mathcal{P}_{n-1}$, where r is a positive real number. The cap is obtained from an arc of the circle S^1_r. The hypersurface $\mathcal{C}_n(r)$ is isometric to the hypersurface defined by

$$\mathcal{C}'_n(r) := \{(x_1, \ldots, x_{n+1}) : x_n^2 + x_{n+1}^2 = r^2\},$$

which can be seen as the image of two functions $f_\pm$:

$$f_\pm(x_1, \ldots, x_n) = \pm\sqrt{r^2 - x_n^2}.$$

This hypersurface is also isometric to the graph $\mathcal{C}''_n(r)$ of the translated functions $r - f_\pm$. Let us define g as

$$g(x_1, \ldots, x_n) := r - \sqrt{r^2 - x_n^2}.$$

We have $g(0, \ldots, 0) = 0$ and $(\partial g/\partial x_i)(0, \ldots, 0) = 0$ for all i. Then, the second fundamental form at the origin of $\mathcal{C}''_n(r)$ reduces to the Hessian matrix of g at the origin whose coefficients are all 0 except the last diagonal one, which is equal to $1/r$. Consequently, the mean curvature of $\mathcal{C}''_n(r)$ at the origin is simply $1/nr$. Thus the mean curvature of $\mathcal{C}_n(r)$ at any of its points is equal to $1/nr$.

The total curvature of the cap approximating the PL-hypersurface is thus equal to the integral over the cap of $1/nr$. Since the cap is tangent to the PL-hypersurface, then, at the contact point, the cap and the PL-hypersurface have the same normal vectors. This means that the angle defining the arc of S^1_r of the cap is equal to π minus the angle Θ between the two normal vectors at the contact points, which are simply the normal vectors of the hyperplanes whose intersection is approximated by the cap. Thus the total mean curvature over the cap is

$$\|\tau_{ij} \cap p\| r(\pi - \Theta_{ij}) \frac{1}{nr}$$

where $\tau_{ij} \cap p$ represents the intersection of τ_{ij} with the neighborhood of p. Now if we suppose that in the neighborhood on the hypersurface around the vertex p the mean curvature H_p is constant, then the total mean curvature over the neighborhood is equal to $H_p\|p\|$. Hence

$$H_p\|p\| = \sum_{\tau_{ij} \in \mathrm{St}^2(p)} \|\tau_{ij} \cap p\| \frac{(\pi - \Theta_{ij})}{n}$$

From this we obtain

$$H_p = \frac{1}{n\|p\|} \sum_{\tau_{ij} \in \mathrm{St}^2(p)} (\pi - \Theta_{ij}) \|\tau_{ij} \cap p\|$$

If we suppose that when computing the mean curvature at all vertices p, the volume neighborhoods around p divide every $(n-1)$-simplex into n subsimplexes of the same volume (*e.g.*, as in a barycentric configuration), then it holds that $\|\tau_{ij} \cap p\| = (1/n)\|\tau_{ij}\|$ and thus:

$$\hat{H}(p) = \frac{1}{n^2\|p\|} \sum_{\tau_{ij} \in \mathrm{St}^2(p)} (\pi - \Theta_{ij}) \|\tau_{ij}\|.$$

5. Experimental evaluation and results

5.1. Analytic surfaces

We evaluate our proposed operator by running our implementation on analytic functions and comparing the result to the corresponding known analytic mean curvature values. The functions are those suggested by Hamann [25] as well as a second trigonometric function. They are as follows:

1. Quadratic polynomial: $0.4(x^2 + y^2 + z^2)$.
2. Quadratic polynomial: $0.4(x^2 - y^2 - z^2)$.
3. Cubic polynomial: $0.15(x^3 + 2x^2y - xz^2 + 2y^2)$.
4. Exponential: $\exp(-0.5(x^2 + y^2 + z^2))$.
5. Trigonometric: $0.1(\cos(\pi x) + \cos(\pi y) + \cos(\pi z))$.
6. Trigonometric: $\sin(\pi x) + \sin(\pi y) + \sin(\pi z)$.

The functions are sampled in the $[-1, 1]^3$ real interval using three different approaches. In the first, we use a uniform grid of samples which is then tessellated using a stencil Voronoi approach. That is to say, a cube is Voronoi-tessellated in the 3D domain and then repeated over the entire grid. In the second approach, we create irregular tessellations. For a given number of vertices, we randomly sample the interior of the interval and Delaunay-tessellate said samples. In order to discourage poorly shaped tetrahedra, we uniformly sample the boundary, enforce a minimum distance constraint during the sampling, and relax the final mesh with 100 iterations of Laplacian smoothing. Finally, we also use diamond meshes [37], a multi-resolution approach that, given an approximation error, can generate a non-uniform mesh of well-shaped tetrahedra that approximates the field. Fig. 1 illustrates these approaches on a 2D slice of the sixth function above.

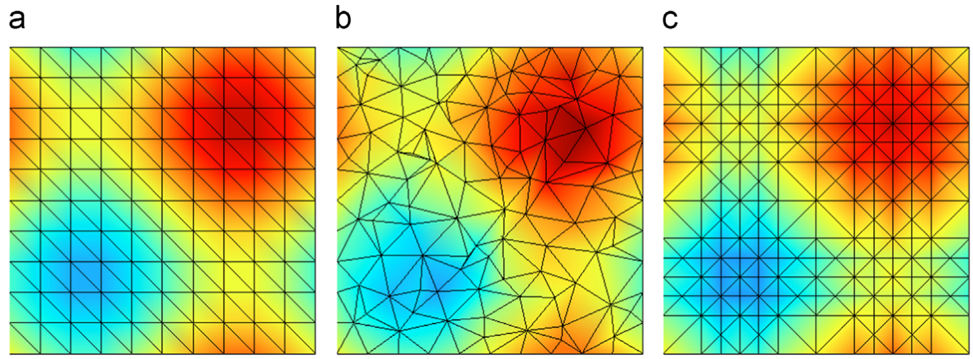

Fig. 1. A $z=0$ slice of the different tessellations used (lowest resolution for each shown for illustrative purposes) illustrated on our sixth analytic function; (a) regular grid, (b) irregularly sampled, (c) diamond mesh.

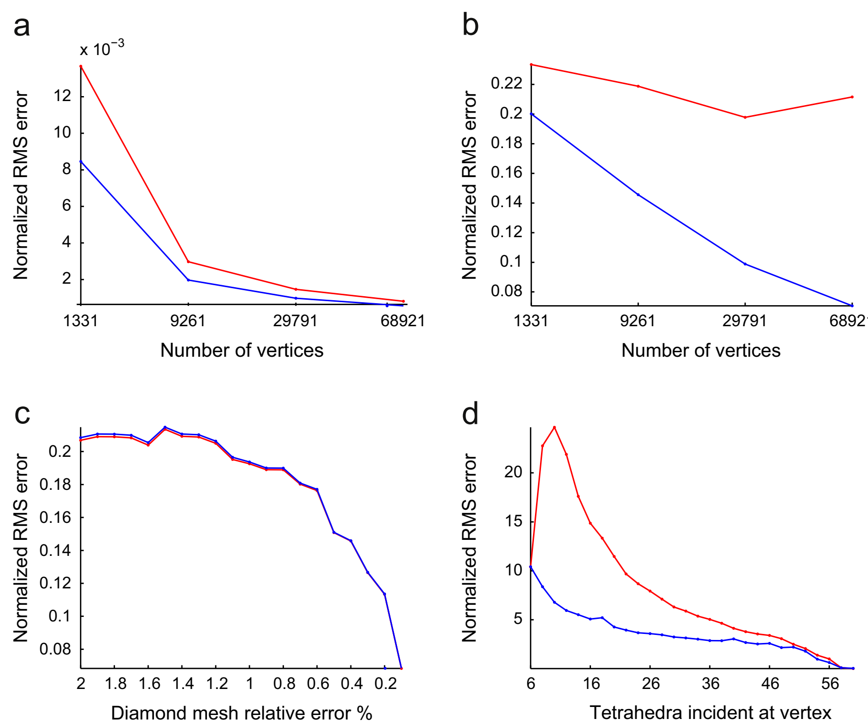

Fig. 2. Average normalized RMS error of weighted distortion (blue) vs the Laplace–Beltrami operator (red) as a function of increasing resolution on (a) uniform grid, (b) irregular, and (c) diamond meshes, and also (d) as a function of the number of tetrahedra incident per vertex on irregular tessellations. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this article.)

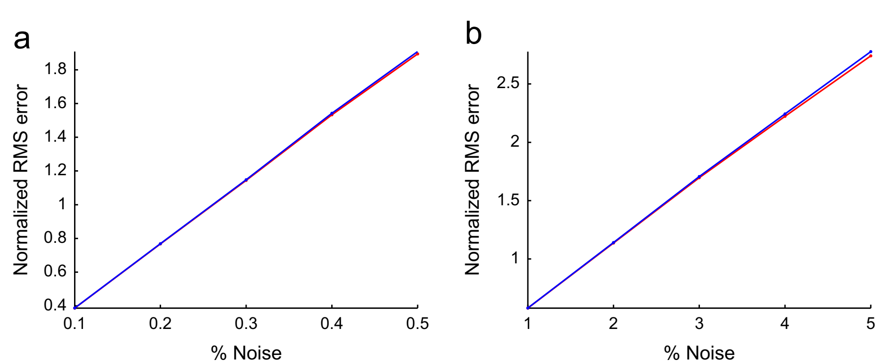

Fig. 3. Average normalized RMS error of weighted distortion (blue) vs the Laplace–Beltrami operator (red) on regular grid tessellations of fixed resolution and increasing Gaussian noise. (a) Noise is added in the vertical direction with standard deviation as a percentage of the field range. (b) Noise is added in the surface normal direction with standard deviation as a percentage of the average edge length. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this article.)

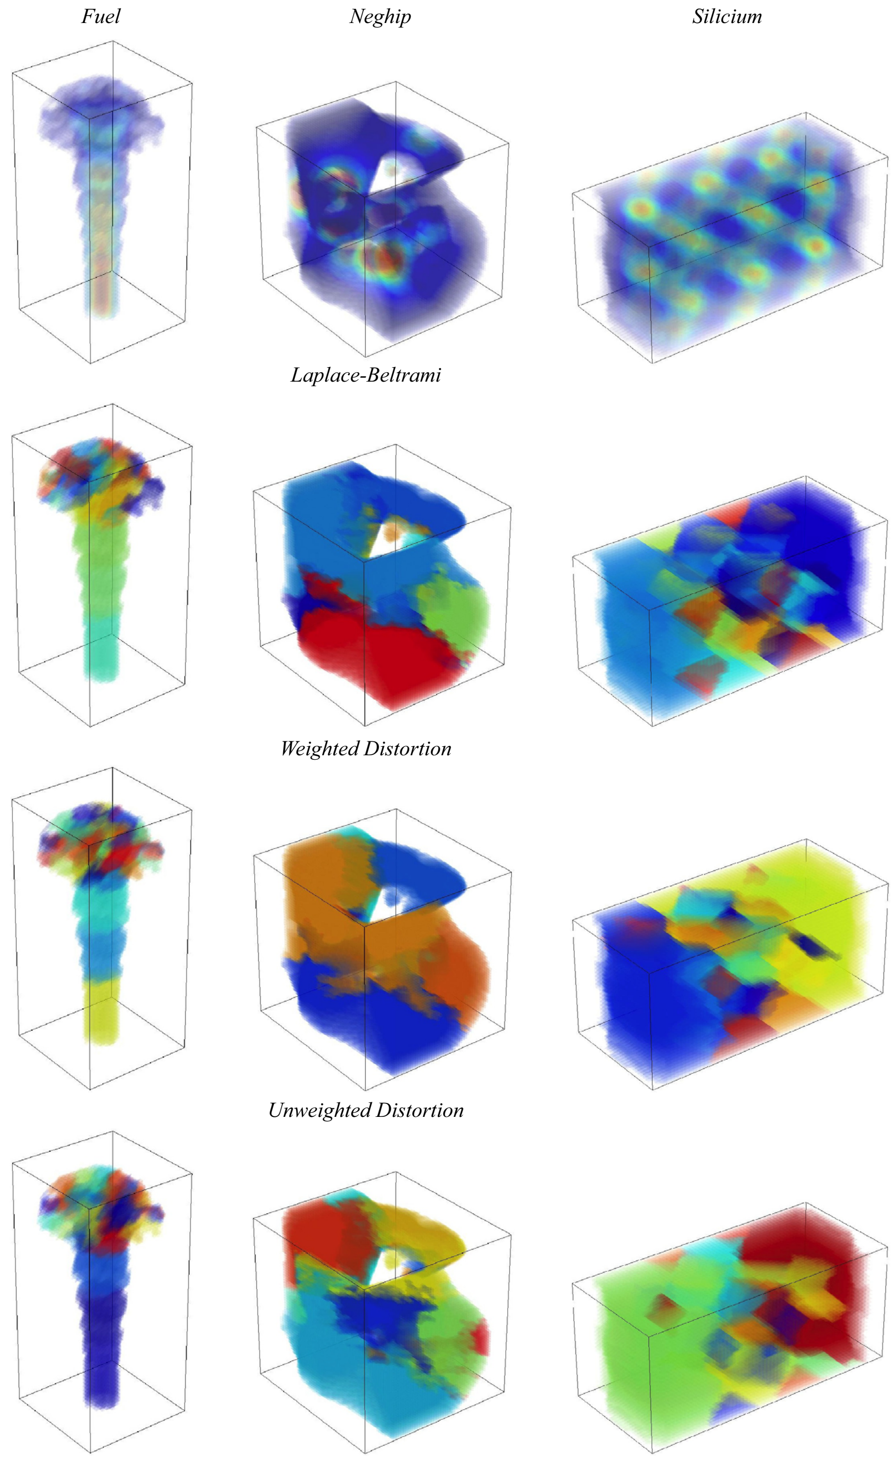

Fig. 4. Segmentations of fuel, neghip, and silicium volume datasets obtained using a hierarchical descending Morse complex on the Laplace–Beltrami, weighted and unweighted distortion fields.

Our first experiments show how the mean curvature error changes with increasing resolution. We use the Root Mean Square (RMS) error normalized by the range (maximum minus minimum) of analytic values taken by each function within the interval and averaged over the six functions. Fig. 2 illustrates these results.

In Fig. 2a, we see the result on uniform grids. While both operators converge, the weighted distortion does so more quickly, achieving an average reduction in error of 29% compared to the Laplace–Beltrami operator. On the non-grid tessellations, both operators are less well-behaved. We remedy this by smoothing the estimates using 50 iterations of local averaging. Fig. 2b and c show these results on the non-grid meshes.

We also compared, in the irregular tessellations, the estimation error as a function of the number of tetrahedra incident on each vertex. For a fixed incidence number, we evaluate the normalized RMS error over the vertices with this valence and average the results over all irregular tessellations and all functions. While both operators converge to the analytic values of mean curvature as

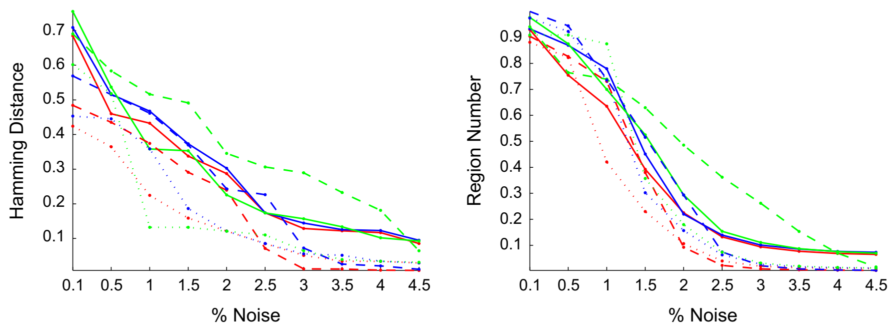

Fig. 5. Comparison of the stability of the segmentations under increasing noise. Left: Hamming distance, right: region number. Red: Laplace–Beltrami, blue: weighted distortion, green: unweighted distortion. Solid: fuel, dashed: neghip, dotted: silicium. Note that, since we are measuring similarity to the noiseless case, higher is better. Also note the rate discontinuity in the *x*-axis. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this article.)

vertex valence increases, we found the weighted distortion error to be much lower than that of the Laplace–Beltrami operator. This is illustrated in Fig. 2d.

Finally, we compare the behavior of the operators under increasing noise. For each uniformly sampled mesh, we add two forms of noise. In the first case, to simulate image noise, we add Gaussian noise to the field component of the vertex coordinates as a proportion of the range of analytic values. In the second, we add Gaussian noise in the normal direction as a proportion of the average edge length. As Fig. 3 illustrates, the operators behave very similarly under these conditions.

5.2. Application to segmentation

To further evaluate our proposed operator, we explore an application to the segmentation of 3-manifolds in 4D space. We extend to 3D the intuition that shape boundaries often perceptually align with concavities, which correspond to regions of negative mean curvature. While this intuition is common and we will use it here, it should be noted that if one scales the function by a factor, the Euclidean curvature values may change in a relatively complicated way. Pottmann and Opitz argue that it may be more natural to use isotropic curvatures [38]. Based on the concavity intuition, we apply a hierarchical Morse decomposition of the mean curvature field defined at the vertices of the tessellated 3-manifold [39]. We use the descending Morse complex, which finds segment centers at locations of high (positive) curvature and boundaries at areas of low (negative) curvature. To counter over-segmentation, the algorithm applies hierarchical region-merging based on the notion of persistence, which relates to the "height" difference between adjacent segments.

In our experiments, we consider three datasets obtained from the Volvis library [40]: *fuel*: a simulation of fuel injection into a combustion chamber; *neghip*: a simulation of the spatial probability distribution of the electrons in a high potential protein molecule; and *silicium*: a simulation of a silicium grid. They and their segmentation results using Laplace–Beltrami, weighted and unweighted distortion, are shown in Fig. 4.

For each dataset, we empirically chose a merging threshold that results in a number of segments between 20 and 40. In each case, we add increasing artificial noise to the field, recompute the Laplace–Beltrami and distortion fields, re-obtain the segmentations using the originally chosen threshold for each set, and measure the similarity to the original segmentation. For this last step, we use the Hamming distance metric proposed by Huang and Dom [41,42]. We also found it interesting to compare the ratio of the number of segments in each segmentation n/m, where $m > n$. While this "region number" metric is much less discriminative, it gives an intuitive sense of how the number of segments grows as a result of noise. Given that these images are originally captured on a regular grid, we simulate image noise by adding Gaussian noise to the field component of the data with standard deviation set as a percentage of the field range. The results are illustrated in Fig. 5, the distortion operator in both its weighted and unweighted forms showing higher similarity to the noiseless segmentation under increasing noise.

6. Concluding remarks

We have examined the previously existing notions of distortion and note that they can be divided into *intrinsic* and *extrinsic* categories depending on whether they are defined using the interior angles or the dihedral angles of the tessellation. We then presented a new discrete operator generalizing the notion of *extrinsic distortion* to *n*D and derived a weighting that can be used to compute mean curvature on such surfaces. We analyzed the behavior of the operator on 3-manifolds in 4D, comparing it to the well known Laplace–Beltrami operator, using ground-truth analytic surfaces with varying conditions of resolution, sampling distribution, and noise. We also investigate it in the context of an application that uses the mean curvature field to obtain a volumetric segmentation, examining the stability of the segmentations under increasing image noise. In each case we showed that extrinsic distortion behaves similarly or better than the Laplace–Beltrami operator while being intuitively simple and easy to implement.

Future work includes increasing the robustness of the operator under conditions of irregular tessellation. Our method could also be applied to the segmentation and analysis of 3-manifold hypersurfaces that are not graphs of 3D scalar fields, as is fully permitted by the current formulation and implementation. Lastly, other applications of mean curvature in higher dimensions are open to investigation, including visualization, registration, matching, alignment, and simplification of volumetric datasets.

Acknowledgments

This work has been partially supported by the National Science Foundation under Grant no. IIS-1116747, and by the Italian Ministry of Education and Research under the PRIN 2009 program.

References

[1] Heckbert PS, Garland M. Optimal triangulation and quadric-based surface simplification. Comput Geometry Theory Appl 1999;14(1–3):49–65.
[2] Sheffer A. Model simplification for meshing using face clustering. Comput Aided Des (CAD) 2001(33):925–34.
[3] Rusinkiewicz S, Levoy M. Efficient variants of the ICP algorithm. In: Proceedings of 3D digital imaging and modeling, 2001. p. 145–52.

[4] Ohtake Y, Belyaev A, Seidel H-P. Ridge-valley lines on meshes via implicit surface fitting. ACM Trans Graph 2004;23(3):609–12.
[5] Rusinkiewicz S. Estimating curvatures and their derivatives on triangle meshes. In: Proceedings of the symposium on 3D data processing, visualization, and transmission, 2004.
[6] Lévy B, Petitjean S, Ray N, Maillot J. Least squares conformal maps for automatic texture atlas generation. In: ACM SIGGRAPH. ACM Press; 2002. p. 362–71.
[7] Lavoue G, Dupont F, Baskurt A. A new cad mesh segmentation method, based on curvature tensor analysis. Comput Aided Des (CAD) 2005;37(10):975–87.
[8] Gatzke T, Grimm C. Estimating curvature on triangular meshes. Int J Shape Model 2006;12(1):1–29.
[9] Mitra NJ, Guibas LJ, Pauly M. Partial and approximate symmetry detection for 3D geometry. In: ACM SIGGRAPH, vol. 25(3), 2006. p. 560–8.
[10] Meyer M, Desbrun M, Schroder M, Barr AH. Discrete differential-geometry operators for triangulated 2-manifolds. In: Hege H-C, Polthier K, editors. Proceedings VisMath 2002, 2003. p. 35–57.
[11] Kalogerakis E, Simari P, Nowrouzezahrai D, Singh K. Robust statistical estimation of curvature on discretized surfaces. In: Proceedings of the symposium on geometry processing, 2007. p. 13–22.
[12] Kalogerakis E, Nowrouzezahrai D, Simari P, Singh K. Extracting lines of curvature from noisy point clouds. Comput Aided Des 2009;41(4):282–92.
[13] Ponchio F, Hormann K. Interactive rendering of dynamic geometry. IEEE Trans Vis Comput Graph 2008;14(4):914–25.
[14] Carno MPD. Differential geometry of curves and surfaces. Englewood Cliffs, NJ: Prentice-Hall Inc.; 1976.
[15] Gatzke T, Grimm C. Estimating curvature on triangular meshes. Int J Shape Model 2006;12:1–29.
[16] Hahmann S, Belayev A, Busé L, Elber G, Mourrain B, Rössl C. Shape interrogation. In: Floriani LD, Spagnuolo M, editors. Shape analysis and structuring. Series on mathematics and visualization. Springer-Verlag; 2009. p. 1–51.
[17] Hamann B. Curvature approximation for triangulated surfaces, vol. 8. Springer-Verlag; 139–53.
[18] Stokely E, Wu S. Surface parametrization and curvature measurement of arbitrary 3-D objects: five practical methods. IEEE Trans Pattern Anal Mach Intell 1992:833–9.
[19] Pottmann H, Wallner J, Huang Q, Yang Y-L. Integral invariants for robust geometry processing. Comput Aided Geometric Des 2009;26:37–60.
[20] Chen X, Schmitt F. Intrinsic surface properties from surface triangulation. In: Proceedings of the European conference on computer vision (ECCV), 1992. p. 739–43.
[21] Taubin G. Estimating the tensor of curvature of a surface from a polyhedral approximation. In: Proceedings ICCV'95, 1995. p. 902–7.
[22] Page DL, Sun Y, Koschan A, Abidi M. Normal vector voting: crease detection and curvature estimation on large, noisy meshes. Graphical Models 64(3/4):199–229.
[23] Langer T, Belyaev AG, Seidel H-P. Exact and interpolatory quadratures for curvature tensor estimation. Comput Aided Geometric Des 2007;5:443–63.
[24] Cohen-Steiner D, Morvan J-M. Restricted Delaunay triangulations and normal cycle. In: Symposium on computational geometry, 2003. p. 312–21.
[25] Hamann B. Curvature approximation of 3D manifolds in 4D space. Comput Aided Geometric Des 1994;11(6):621–33.
[26] Mesmoudi MM, De Floriani L, Port U. Discrete distortion in triangulated 3-manifolds. Comput Graph Forum 2008;27(5):1333–40.
[27] Mesmoudi MM, De Floriani L, Magillo P. Morphology analysis of 3D scalar fields based on Morse theory and discrete distortion. In: ACM international symposium on advances in geographic information systems, Seattle, 2009.
[28] Weiss K, De Floriani L, Mesmoudi MM. Multiresolution analysis of 3D images based on discrete distortion. In: Proceedings of the international conference on pattern recognition, 2010. p. 4093–6.
[29] De Floriani L, Iuricich F, Magillo P, Mesmoudi MM, Weiss K. Discrete distortion for 3D data analysis. Vis Med Life Sci 2012;II:3–25.
[30] Aleksandrov P. Topologia Combinatoria. Torino: Edizioni Scientifiche Einaudi; 1957.
[31] Troyanov M. Les surfaces Euclidiennes à singularités coniques. L'Enseignement Mathématique 1986;32:79–94.
[32] Mesmoudi MM, De Floriani L, Magillo P. Concentrated Curvature for Mean Curvature Estimation in Triangulated Surfaces. In: CTIC, 2012. p. 79–87.
[33] Jin M, Kim J, Gu X. Discrete surface Ricci flow: theory and applications. In: Proceedings 12th IMA international conference on mathematics of surfaces XII. Lecture notes in computer science, vol. 4647, 2007. p. 209–32.
[34] Xiaotian Y, Miao J, Feng L, Xianfeng G. Discrete curvature flow for hyperbolic 3-manifolds with complete geodesic boundaries. In: Proceedings Eurographics symposium on geometry processing. Computer graphics forum, vol. 27, 2008.
[35] Dyn N, Hormann K, Sun-Jeong K, Levin D. Optimizing 3D triangulations using discrete curvature analysis. In: Mathematical methods for curves and surfaces: Oslo 2000. Vanderbilt University; 2001. p. 135–46.
[36] Gatzke TD, Grimm CM. Estimating curvature on triangular meshes. Int J Shape Model 2006;12(1):1–28.
[37] Weiss K, De Floriani L. Simplex and diamond hierarchies: models and applications. Comput Graph Forum 2011;30(8):2127–55.
[38] Pottmann H, Opitz K. Curvature analysis and visualization for functions defined on Euclidean spaces or surfaces. Comput Aided Geometric Des 1994;11:655–74.
[39] Čomić L, De Floriani L, Iuricich F. Dimension-independent multi-resolution Morse complexes. Comput Graph 2012;36(5):541–7.
[40] Volvis library ⟨http://www.volvis.org/⟩.
[41] Huang Q, Dom B. Quantitative methods of evaluating image segmentation. In: International conference on image processing, 1995. Proceedings, vol. 3, October. p. 53–6.
[42] Chen X, Golovinskiy A, Funkhouser T. A benchmark for 3D mesh segmentation. In: Proceedings of ACM SIGGRAPH. ACM; 2009. p. 73.