

As-Conformal-As-Possible Discrete Volumetric Mapping

Gilles-Philippe Paillé, Pierre Poulin

LIGUM, Dept. I.R.O., Université de Montréal

Abstract

In this paper, we tackle the problem of generalizing conformal maps to volumetric meshes. Current methods seek for harmonicity but unfortunately, no computational methods optimize conformality in the volumetric context. As it is proven that conformal maps do not exist for general volume transformations, we seek to optimize shape preservation with a generalization of the Cauchy-Riemann equations. Our algorithm is fast and easily adaptable to existing harmonic mapping methods. Compared to harmonic maps, results show improvements on both angular and volumetric energy measures at a cost below 1% of total computations. The method extends well in any dimension and several research areas could benefit from our derivations of volumetric conformal optimization.

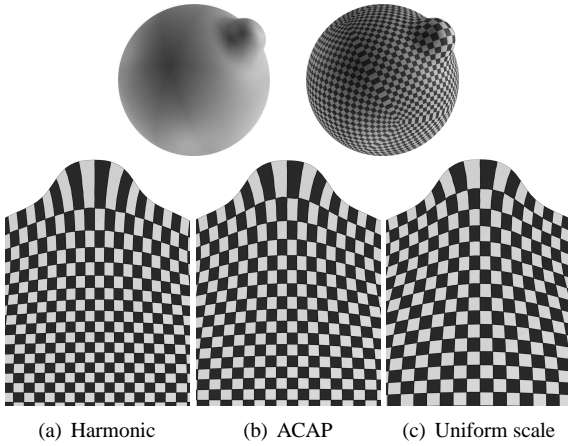


Figure 1: An identical planar cut through a sphere with a small bump that is parameterized with (a) the Laplace operator, (b) the our operator, and (c) the uniform scale operator ($\omega = 0.6125$).

1. Introduction

Parameterization plays an important role in computer graphics. It is best known for texturing 3D models but it is also used for shape matching and analysis, remeshing, spline construction, physical simulation, etc. [1]. While being mostly studied for surfaces, volumetric parameterization has recently shown an increased interest, notably in the domain of engineering and medical imaging [2], where structural evaluations require a good distribution of internal volume elements.

Most of the current parameterization techniques are based on harmonic maps that are obtained by solving the Laplace equation. They offer good deformation properties for the interior of the object, but the boundary often suffers from higher nonuniform scaling, as shown in Figure 1.

This paper revisits the Laplace operator by generalizing the well known conformal maps to the volumetric context, thus leading to volumetric as-conformal-as-possible (ACAP) maps.

Conformality has important qualities such as angle and local shape preservation, and thus it is natural to ask for similar properties when extending from surfaces to volumes. As far as we know, no such attempts to generalize conformal maps to the 3D context has been made so far.

By generalizing the Cauchy-Riemann equations, we obtain a first order matricial operator that optimizes conformality in each of the three canonical orthogonal planes. Minimizing the error in the least-squares sense leads to a simple second order operator similar to the Laplace operator. In fact, a family of operators can be obtained using the same derivations, leading to other forms of energy minimization, thus showing the generality of the equations. Indeed, we show that orthogonality and uniform scaling can be weighted to obtain the desired map constraints.

The mathematical expression being similar to the Laplace operator, many parameterization methods could benefit from our approach. We also believe that this could have several applications in areas where the Laplace operator is currently used, such as in physically-based and non-physically-based animations [3].

The paper is organized as follows. Section 2 reviews the essential surface and volumetric parameterization methods, and then introduces the contributions of our method. Section 3 describes the energy minimization technique, starting from the continuous case to the discrete case, before synthesizing the whole method in a simple algorithm. Numerical and visual results are provided in Section 4, and conclusions in Section 5.

2. Related Work

Surface parameterization has been extensively studied in the past. Even though volumetric parameterization is built on top of surface parameterization, reviewing the latter is beyond the scope of this paper. For good surveys on these techniques, see

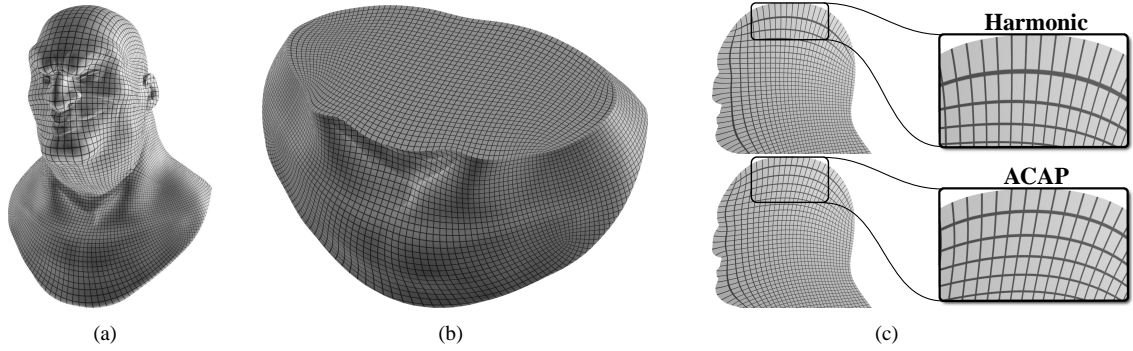


Figure 2: As-conformal-as-possible volumetric map of (a) the Mudface model to a polycube along with (b) a iso-v cut and (c) a side iso-u cut of the mesh. The iso-u cut shows differences between harmonic mapping and our mapping. Note the increased uniformity near the surface borders.

[4, 5, 6, 1]. In our implementation, we use a simple piecewise-linear parameterization approach that optimizes conformality on arbitrary domains. Other methods could be used as well.

Volumetric parameterization techniques began to appear in 2003. Wang et al. [7] introduce a simple method for computing volumetric harmonic maps by conformally mapping the surface of a genus zero object to a sphere, and then optimizing for the interior using methods similar to those used to solve the heat equation. The technique is also applied to brain mapping [8]. While the method is not strictly restricted to genus zero objects, the authors do not describe how to extend it to domains of arbitrary topology.

Many papers have emerged since 2007. The method of fundamental solutions (MFS) uses the electric charge metaphor to compute harmonic maps between shapes sharing the same topology [9, 10]. This method is improved by aligning features on the codomain shape [11]. The MFS is also used to compute harmonic maps on star-shaped domains [12, 13].

Xia et al. [14] map the object to a polycube [15] decomposed into a floor, walls, and a ceiling. A harmonic scalar function is computed on the object and the polycube with boundary values specified on the floor and the ceiling. The final map is computed by following the integral curves of the gradient of the harmonic function on both domains.

A trivariate B-spline can be fitted to an object by mapping it harmonically to the cylinder [2]. The method can be generalized to other trivariate representations with more complex topologies by segmenting the volume into trivial parts, ensuring coherence between volumetric patches [16]. A generalized polycube representation can manage non-trivial topologies such as volumetric Möbius bands and identified cube borders [17].

Another method consists in mapping the object to a polycube by ensuring orthogonality on the first vertex layer, and then minimizing an edge-length constraint on interior vertices [18]. It can lead to sheared interior elements in high curvature areas. The method is used to simulate diffusion of light in heterogeneous translucent media.

Yin et al. [19] map a volume with complex topologies to the periodic hyperbolic space using a discrete curvature flow. This method generalizes to 3D the fact that high genus surfaces can be mapped to the hyperbolic plane without singularities.

Hexahedral meshing methods share similar concepts with

volumetric parameterization, and three recent techniques are worth noticing. A mapping to the polycube is computed using rotation- and position-driven deformations [20]. The resulting parameterization is used to map the hexahedra of the polycube to the original object in order to produce a hexahedral mesh. The CubeCover method [21] uses a technique similar to QuadCover [22], where each element is optimized to be aligned to a given vector field. Using an anisotropic metric, Lévy and Liu [23] produce a hex-dominant mesh using Centroidal Voronoi Tesselation.

Thin-shell parameterization considers that the surface of the mesh is a volumetric shell that needs to be taken into account in the mapping computation. The shell can be mapped to a polycube [24] or unfolded to a periodic plane of constant thickness [25]. Both papers use harmonic maps. Li et al. [26] use a harmonic field to map a thin shell representation of a standard hexahedral facial model to a patient-specific facial model while maximizing orthogonality. Although the method shares some properties with conformality, uniform scaling is not part of the optimization.

Contributions. This paper presents the first method that introduces a simple and practical definition of conformal optimality for volumes. The generalization of the Cauchy-Riemann equations in conjunction with the generalized Dirichlet energy leads to a fast discrete least-squares approximation of conformality in 3D. The result is fundamental, theoretically founded, and can be incorporated in many existing methods without difficulty. The linearity of the operator plays an important role in making the method fast, straightforward, and simple.

As shown in Section 4, our method performs better than harmonic maps for both angle and volume energy measures.

3. As-Conformal-As-Possible Mapping

3.1. Continuous Case

In 2D, conformal maps are functions that locally preserve the shape of the original space, i.e., the Jacobian matrix at each point of the space is a similarity transformation. For function $\mathbf{f} : \mathcal{S}_2 \rightarrow \mathbb{R}^2$ defined over surface \mathcal{S}_2 , this constraint is expressed by the Cauchy-Riemann equations

$$\begin{bmatrix} \partial_x & -\partial_y \\ \partial_y & \partial_x \end{bmatrix} \begin{bmatrix} f_x \\ f_y \end{bmatrix} = 0.$$

In 3D, the local shape preservation constraint is too strong to be perfectly respected [27], but it can be as optimal as possible. For this reason, we propose to apply the Cauchy-Riemann equations on each of the three canonical orthogonal planes as a way to maximize local rigidity. For function $\mathbf{f} : \mathcal{S}_3 \rightarrow \mathbb{R}^3$ defined over volumetric shape \mathcal{S}_3 , we define the new constraint $\mathcal{D}\mathbf{f} = 0$ where

$$\mathcal{D} = \begin{bmatrix} 0 & \partial_y & -\partial_z \\ 0 & \partial_z & \partial_y \\ \partial_x & 0 & -\partial_z \\ \partial_z & 0 & \partial_x \\ \partial_x & -\partial_y & 0 \\ \partial_y & \partial_x & 0 \end{bmatrix}.$$

The operator \mathcal{D} is unfortunately not rotational invariant because each equation considers that its corresponding plane in the parametric domain will not change its orientation during the transformation. We therefore propose to use a technique issued from mesh deformation, namely Warped Stiffness [28], which consists in removing the local rotation of the map at each point before applying the operator. Let \mathbf{R} be the local rotation at point \mathbf{x} , the constraint becomes

$$\mathcal{D}\mathbf{R}^T \mathbf{f} = 0.$$

In this form, the Jacobian matrix of the map is overconstrained. That is why we propose to minimize it in the least-squares sense. To achieve this, we seek to minimize the following energy functional

$$\mathcal{E} = \int_{\mathcal{S}_3} \|\mathcal{D}\mathbf{R}^T \mathbf{f}\|^2 dV.$$

3.2. Discrete Case

We first proceed by discretizing the space and functions defined on this space. Let \mathcal{M} be a tetrahedral mesh with vertices \mathcal{V} and tetrahedra \mathcal{T} . A function is defined over \mathcal{M} by associating a value to every vertex and linearly interpolating inside a tetrahedron. Let $t \in \mathcal{T}$ be a tetrahedron of volume V_t with function value \mathbf{f}_i at vertex v_i . The barycentric interpolation on tetrahedron t is

$$\mathbf{f}(\mathbf{x}) = -\frac{1}{3V_t} \sum_{i=1}^4 (\mathbf{x} \cdot \mathbf{n}_i) \mathbf{f}_i \quad (1)$$

where \mathbf{n}_i is the outside normal of the opposite face of the i th vertex in the tetrahedron, with normal length equal to the area of the face. Notations are illustrated in Figure 3.

The constraints being applied on the values of the Jacobian matrix, we compute the latter by differentiating Equation (1). Using the tensor product [29], we obtain

$$\mathbf{J} = -\frac{1}{3V_t} \sum_{i=1}^4 \mathbf{f}_i \otimes \mathbf{n}_i.$$

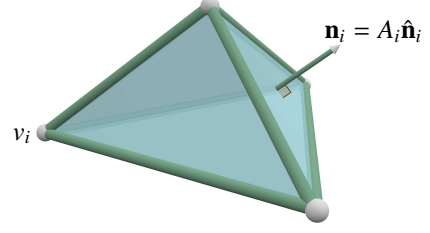


Figure 3: Tetrahedron notation: A_i is the area of the face opposite to vertex v_i , $\hat{\mathbf{n}}_i$ is the unit normal of the face, and \mathbf{n}_i is the area-weighted normal.

Having computed all the derivatives, it is now easy to define our discrete operator D over the tetrahedron as

$$D\mathbf{f} = -\frac{1}{3V_t} \sum_{i=1}^4 \mathbf{D}_i \mathbf{f}_i$$

where \mathbf{D}_i is defined as

$$\mathbf{D}_i = \begin{bmatrix} 0 & n_{i,y} & -n_{i,z} \\ 0 & n_{i,z} & n_{i,y} \\ n_{i,x} & 0 & -n_{i,z} \\ n_{i,z} & 0 & n_{i,x} \\ n_{i,x} & -n_{i,y} & 0 \\ n_{i,y} & n_{i,x} & 0 \end{bmatrix}. \quad (2)$$

At this point, we use a technique similar to Wang et al. [7] and Lévy et al. [30] to derive the least-squares formulation. The goal is to minimize an energy functional over the whole tetrahedral mesh. The discretized map being linear by parts, we define the energy on each individual element. For tetrahedron $t \in \mathcal{T}$ with precomputed local rotation \mathbf{R}_t , the energy functional is

$$E_t = \frac{1}{2} \int_t \|\mathcal{D}\mathbf{R}_t^T \mathbf{f}\|^2 dV_t.$$

To obtain a linear system, we differentiate the discrete energy at vertex $v_i \in \mathcal{V}$ with one-ring vertex neighborhood $\mathcal{N}(v_i)$, obtaining

$$\sum_{j \in \mathcal{N}(v_i) \cup \{i\}} \mathbf{K}_{ij} \mathbf{f}_j = 0 \quad (3)$$

where \mathbf{K}_{ij} is the matrix defined by

$$\mathbf{K}_{ij} = \sum_{t \in \mathcal{T}(v_i, v_j)} \frac{\mathbf{R}_t \mathbf{D}_{t,i}^T \mathbf{D}_{t,j} \mathbf{R}_t^T}{9V_t}$$

where $\mathcal{T}(v_i, v_j)$ are the tetrahedra sharing edge $\{v_i, v_j\}$ and $\mathcal{T}(v_i, v_i)$ are the tetrahedra sharing vertex v_i .

3.3. Algorithm

To solve the system defined by Equation (3) and obtain the as-conformal-as-possible volumetric map, we use standard solvers at each step of the technique. Here is a description of each step.

Step 1 - Surface Map. The first step is to map the surface of the mesh to the surface of the codomain using any method. Note that the quality of the surface map affects the quality of the volumetric map, e.g., singularities and triangle inversions could be propagated to a certain distance in the interior. As long as the surface mapping is appropriate, whatever the mapping used, good results should be expected.

We chose the codomain to be a cube, a sphere, or a polycube [15]. This choice depends on how well the surface maps to each codomain. Except for the simplest objects, we chose polycubes as the preferred codomains. Figure 7 shows typical polycubes for some of the objects tested. The surface map is set to be conformal. We compute the surface map with a Gradient Descent method with adaptative step size. The number of iterations is approximately the number of surface vertices. The threshold is set to 10^{-5} using the average squared residual.

Step 2 - Volumetric Harmonic Map. The second step relies on work by Wang et al. [7] to compute the first approximation of the volumetric part of the map. Using the equations of Section 3.2, we obtain the linear system for a harmonic map by modifying the $\mathbf{K}_{i,j}$ matrices. We define the new coefficients of Equation (3) to

$$\tilde{\mathbf{K}}_{i,j} = \sum_{t \in \mathcal{T}(v_i, v_j)} \frac{\mathbf{n}_{t,i} \cdot \mathbf{n}_{t,j}}{9V_t}.$$

We then obtain a symmetric linear system that can be solved with a Conjugate Gradient method with a Jacobi preconditioner, keeping surface vertices constant. The number of iterations is usually around 150 for each mesh with a threshold of 10^{-8} .

Step 3 - Local Rotations and Coefficients. Using the harmonic map as a first approximation, we can compute the local rotation of each vertex. Using those rotations, we compute the \mathbf{K}_{ij} coefficients.

Step 4 - As-Conformal-As-Possible Map. The last step is to compute the volumetric map with our operator by solving Equation (3) using a Conjugate Gradient method with a Jacobi preconditioner. The number of iterations is also usually around 150 for each mesh with a threshold of 10^{-8} .

We found that iteratively updating local rotations and the mapping only leads to epsilon improvements, which is why we only evaluate local rotations once.

4. Results

This section presents the results of the volumetric maps obtained using our metric. We also compare with alternative metrics and discuss the main differences between these metrics. Note that our aim is to introduce a fast and simple metric that optimizes conformality. For this reason, we focus our attention on the behavior of each metric instead of the final map as a whole.

The method is implemented in C++ in the *Graphite* software [31] and all tests were conducted on an Intel Xenon E5520

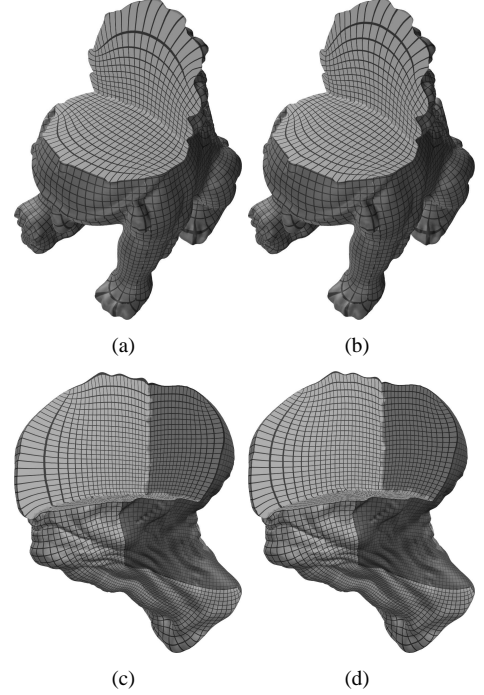


Figure 4: Results of (left) harmonic maps and (right) as-conformal-as-possible maps.

processor without any GPU acceleration technique. Computation times of the interior (i.e., without surface map) are shown in Table 1.

The volumetric parameterization times are approximately multiplied by a factor between three and four using our operator compared to the Laplace operator alone. It is important to note that in general, most of the complete parameterization process is completely absorbed by the surface parameterization as it uses a Gradient Descent method, which is known to be slow. A multi-grid method could be used for this part, but this is left for future work. In fact, the volumetric parameterization times represent generally much less than 1% of that of the surface, which thus far outweighs the increase in time of performing as-conformal-as-possible mapping. Considering the minimal cost and the fact that results are always at least somewhat better, it is worth taking the extra step toward conformality.

From a qualitative point of view, we can note some visual indicators that show improvements. While harmonic maps have very good behavior in the center of the object, quality at the borders is often poor. It is shown in the bumpy sphere of Figure 1 and it is also noticeable in Figure 4. As a compromise, improvements at the borders come at the price of slightly reducing the quality of the interior, but nonetheless the overall quality is improved and the interior has still a good behavior.

To quantitatively measure improvements of the mapping, we use an angle and a volume energy measure based on equations found in the book by Botsch et al. [3] and summarized by Solomon et al. [32]. Adapted to volumes, we obtain

Mesh	$ \mathcal{V} $	$ \mathcal{T} $	Codomain	Harmonic			ACAP		
				$Time$	E_{angle}	E_{volume}	$Time$	E_{angle}	E_{volume}
Sphere	10k	53k	Box	0.18	2.246	2.096	0.89	2.226	2.060
Torus	15k	59k	Polycube	0.12	2.424	2.392	0.58	2.400	2.331
Squirrel	18k	73k	Polycube	0.15	3.118	3.453	0.84	2.638	3.113
Bust	18k	77k	Polycube	0.19	2.829	3.526	0.87	2.660	3.447
Foot	20k	89k	Polycube	0.19	2.441	2.455	0.98	2.372	2.384
Pensatore	30k	150k	Box	0.47	2.554	2.488	2.03	2.441	2.288
Lion Vase	40k	163k	Polycube	0.44	3.324	3.229	2.09	2.869	3.216
Grog	46k	197k	Polycube	0.51	3.521	3.600	2.30	3.174	3.461
Bumpy Sphere	66k	307k	Sphere	2.08	2.003	2.011	5.45	2.003	2.007
			Box	2.01	2.283	2.126	6.53	2.252	2.077
Alien Shark	185k	850k	Polycube	5.16	2.697	2.783	17.48	2.538	2.568
Human Elder	188k	877k	Polycube	5.68	2.598	2.417	30.59	2.524	2.298
Bimba Con Nastrino	203k	969k	Polycube	6.36	2.825	5.912	23.14	2.486	5.443
Chinese Lion	273k	1203k	Polycube	9.37	2.964	3.105	34.71	2.636	3.079
Kitten	342k	1968k	Polycube	23.72	2.668	4.073	74.57	2.530	3.655
Human Torso	389k	1863k	Polycube	16.33	3.056	3.910	53.91	2.571	3.569
Mudface	1183k	5569k	Polycube	48.12	2.634	3.178	154.52	2.469	2.802



Table 1: Quantitative results for different meshes. Timings are in seconds and the optimal minimal value for both energy measures is 2.

$$E_{angle} = \frac{1}{V} \sum_{t \in \mathcal{T}} V_t \left(\frac{\sigma_{t,min}}{\sigma_{t,max}} + \frac{\sigma_{t,max}}{\sigma_{t,min}} \right)$$

$$E_{volume} = \frac{1}{V} \sum_{t \in \mathcal{T}} V_t \left(\sigma_{t,1} \sigma_{t,2} \sigma_{t,3} + \frac{1}{\sigma_{t,1} \sigma_{t,2} \sigma_{t,3}} \right)$$

where V is the total volume of the mesh, $\sigma_{t,i}$ are the singular values of the Jacobian matrix of the tetrahedra transformation, and $\sigma_{t,min}$ and $\sigma_{t,max}$ are respectively the minimum and maximum singular values.

These measures can be numerically unstable near singularities. For this reason, we decided that energy values higher than a fixed threshold should be ignored, due to the bias that they introduce. Even though ignoring values also introduces a bias, the result is still much more representative of the total energy than incorporating unstable data. These extremal values represent generally less than 0.3% of the total volume, which renders the bias negligible when computing global map energies.

From these measures, we can see an improvement on both angle and volume global energies for all tested meshes, as summarized in Table 1. Although not shown in the table, energy variance is also significantly decreased on both measures. Figure 5 shows a color-coded energy distribution of a planar cut through the Chinese Lion model along with a histogram of this distribution for the entire tetrahedral mesh.

Note that bijectivity is not guaranteed for large deformations. In fact, linear constraints do not generally lead to barrier metrics, which means that applying enough force on an element

will eventually invert it. For this reason, codomains should be chosen carefully. While polycubes in this paper are built manually, Gregson et al. [20] provide an automatic geometry-friendly polycube construction method that could be used to minimize element inversion.

4.1. Alternative Metrics

Other metrics can also optimize conformality in higher dimensions. Because the Jacobian matrix of a conformal map should be a scaled rotation, the following energy functional could be used

$$E_t = \frac{1}{2} \int_t \|\mathbf{J} - s\mathbf{R}\|^2 dV_t$$

where $s = |\mathbf{J}|^{\frac{1}{3}}$.

We found that the results with this metric are better than harmonic maps, but not better than ACAP maps, as shown in Table 2. An average of four iterations are needed to converge for tested meshes.

The metric developed by Gregson et al. [20] could also be considered

$$\tilde{v}_i - \frac{1}{N} \sum_{j=1}^N \tilde{v}_j = \frac{1}{N} \sum_{j=1}^N \left(\frac{\mathbf{R}_i + \mathbf{R}_j}{2} \right) \cdot (v_i - v_j)$$

where $N = |\mathcal{N}(v_i)|$ and \tilde{v}_i is the mapped vertex position.

Unfortunately, this operator is highly dependent on the local topology of the vertex and less on the geometry of the volume, in the same vein as mass-spring systems.

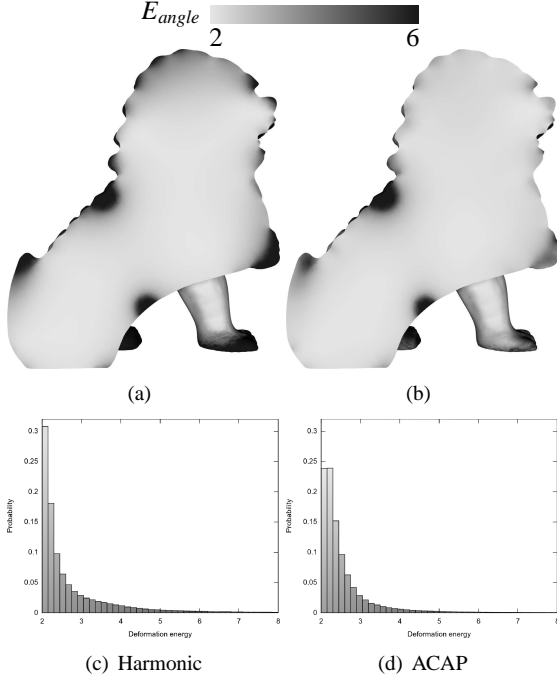


Figure 5: Angle energy measure for the Chinese Lion mapped to a polycube. A planar cut with color-coded energy and the associated deformation distribution for (left) the harmonic map and (right) the as-conformal-as-possible map. Note that 86% of the volume has an energy value below 3 for our map compared to 76% for the harmonic map.

Li et al. [26] developed a metric trying to improve orthogonality, but it has a preferred direction. While it is appropriate for their application, it is an unwanted behavior for general volumetric mappings.

Cage-based deformations could also be considered as potential mapping metrics. The surface map acting as the cage deformation, interior points are moved accordingly to the formulation of the method. Green coordinates [27], being not interpolatory, lead to highly distorted elements near boundaries since vertices can stand outside the cage after deformation. Harmonic maps being uniquely defined by their boundary values, Harmonic coordinates [33] were not considered here. Finally, Mean value coordinates (MVC) [34] gave slightly better results than harmonic maps regarding energy measures on the tested meshes, as shown in Table 2. However, the computation time is prohibitive for large meshes and the negative value property of MVC can lead to high distortion in concave parts.

	SRJ	[GSZ11]	[LLT11]	MVC
Sphere	2.242	2.258	2.234	2.245
Pensatore	2.493	2.560	2.638	2.442
Alien Shark	2.637	2.584	2.746	2.585

Table 2: Angle energy measures of alternative metrics for three meshes. Similar results were obtained with all tested meshes and the volume measure E_{volume} . Compared metrics are the scaled rotated Jacobian matrix (SRJ), the metric from Gregson et al. [20], the metric from Li et al. [26], and Mean value coordinates [34].

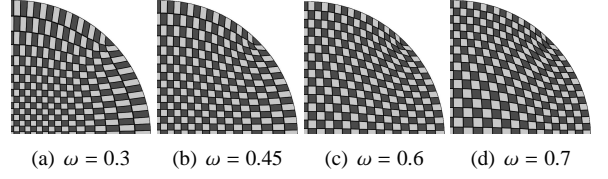


Figure 6: Uniform scale optimization for the sphere parameterized to a box.

4.2. Uniform Scaling

We can also introduce a parameter to operator \mathcal{D} to add a degree of freedom to the results. A logical parameterization would be to separate the uniform scaling constraint from the orthogonality constraint using a weighting factor $\omega \in [0, 1]$, noting $\bar{\omega} = 1 - \omega$. Explaining only for the discrete case, we redefine the matrix of Equation (2) to be

$$\mathbf{D}_{i,\omega} = 2 \begin{bmatrix} 0 & \omega n_{i,y} & -\omega n_{i,z} \\ 0 & \bar{\omega} n_{i,z} & \bar{\omega} n_{i,y} \\ \omega n_{i,x} & 0 & -\omega n_{i,z} \\ \bar{\omega} n_{i,z} & 0 & \bar{\omega} n_{i,x} \\ \omega n_{i,x} & -\omega n_{i,y} & 0 \\ \bar{\omega} n_{i,y} & \bar{\omega} n_{i,x} & 0 \end{bmatrix}$$

while all other equations remain the same.

From experiments, we found that $\omega \in [0.3, 0.7]$ gives best results, $\omega = 0.5$ being the original operator. Extremal values of ω must be avoided because singularities of the surface map may cause degeneracies in the volumetric map. Figure 6 shows mappings computed for different values of ω .

5. Conclusion and Future Work

We introduced an operator that minimizes a 3D conformal energy similar to the Cauchy-Riemann equations in 2D. We showed that this operator can be parameterized such that uniform scaling and orthogonality constraints can be weighted as desired. In fact, any system of constraints using only first-order derivatives can be used with the same derivations. The generalization in any dimension is also trivial.

The derivations are based on a mathematically sound approach. Although the results might not always visually appear significantly improved, they are consistently better than not using our metric, and at a cost of less than 1% of the total computations. Like most other methods, we rely on a conformal surface map. Our polycube map helps to obtain better results, but as shown in Figure 7, simple polycubes prove sufficient to achieve reasonable surface map.

Several methods could benefit from our work. Indeed, most papers that use volumetric harmonic maps as a central technique can be adapted without much difficulty to our maps. This is due to the fact that the core concepts of the mapping remain unchanged, while adding more generality and flexibility to these methods.

Our operator shares the same limitations than the Laplace operator, i.e., that the one-to-one mapping and the convergence are not guaranteed for meshes that lead to ill-conditioned linear systems.

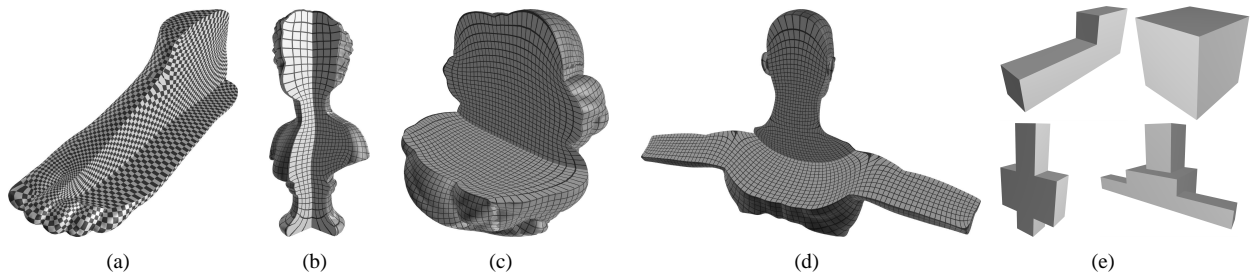


Figure 7: As-conformal-as-possible maps of (a) Foot, (b) Bust, (c) Pensatore, and (d) Human Torso. (e) Codomains of each object.

In the future, we will look for a rotational-invariant operator, hopefully linear, thus eliminating the need for harmonic map pre-processing. We also aim at letting the boundary evolve on the codomain surface as the interior is optimized. Such method would give more room for optimization and could lead to a much better energy minimization. However, this generalization is not straightforward and would need nonlinear optimization.

As parameterization can be seen as a deformation of a shape, we believe that the presented method can have several applications in animation, notably for cage-based, skeleton-based, and point-based deformations.

Acknowledgements

The models are courtesy of Nick Zuccarello (Alien Shark, Human Elder, Mudface) and Aim@Shape shape repository.

References

- [1] K. Hormann, K. Polthier, A. Sheffer, Mesh parameterization: Theory and practice, in: SIGGRAPH Asia Course Notes, 2008.
- [2] T. Martin, E. Cohen, R. M. Kirby, Volumetric parameterization and trivariate B-spline fitting using harmonic functions, *Computer Aided Geometric Design* 26 (6) (2009) 648–664.
- [3] M. Botsch, L. Kobbelt, M. Pauly, P. Alliez, B. Lévy, *Polygon Mesh Processing*, AK Peters, 2010.
- [4] M. S. Floater, K. Hormann, Surface parameterization: a tutorial and survey, in: *Advances in Multiresolution for Geometric Modeling*, Springer, 2005, pp. 157–186.
- [5] A. Sheffer, E. Praun, K. Rose, Mesh parameterization methods and their applications, *Found. Trends. Comput. Graph. Vis.* 2 (2006) 105–171.
- [6] K. Hormann, B. Lévy, A. Sheffer, Mesh parameterization: Theory and practice, in: *ACM SIGGRAPH Course Notes*, 2007.
- [7] Y. Wang, X. D. Gu, S.-T. Yau, Volumetric harmonic map, *Communications in Information and Systems* 3 (3) (2003) 192–202.
- [8] Y. Wang, X. D. Gu, T. Chan, P. Thompson, S.-T. Yau, Volumetric harmonic brain mapping, in: *IEEE International Symposium on Biomedical Imaging: Nano to Macro*, 2004, Vol. 2, 2004, pp. 1275–1278.
- [9] X. Li, X. Guo, H. Wang, Y. He, X. D. Gu, H. Qin, Harmonic volumetric mapping for solid modeling applications, in: *ACM Symposium on Solid and Physical Modeling*, SPM '07, 2007, pp. 109–120.
- [10] X. Li, X. Guo, H. Wang, Y. He, X. D. Gu, H. Qin, Meshless harmonic volumetric mapping using fundamental solution methods, *IEEE Transactions on Automation Science and Engineering* 6 (3) (2009) 409–422.
- [11] X. Li, H. Xu, S. Wan, Z. Yin, W. Yu, Feature-aligned harmonic volumetric mapping using MFS, *Computers & Graphics* 34 (3) (2010) 242–251, *Shape Modeling International (SMI) Conf.* 2010.
- [12] X. Li, Nondegeneracy of harmonic volumetric parameterization on star-shaped domains, Tech. rep., Electrical and Computer Engineering Department, Louisiana State University (2010).
- [13] J. Xia, Y. He, S. Han, C.-W. Fu, F. Luo, X. D. Gu, Parameterization of star-shaped volumes using Green's functions, in: *Advances in Geometric Modeling and Processing*, Vol. 6130 of *Lecture Notes in Comp. Science*, Springer, 2010, pp. 219–235.
- [14] J. Xia, Y. He, X. Yin, S. Han, X. D. Gu, Direct-product volumetric parameterization of handlebodies via harmonic fields, in: *Shape Modeling International Conf. (SMI)*, 2010, pp. 3–12.
- [15] M. Tarini, K. Hormann, P. Cignoni, C. Montani, Polycube-maps, in: *ACM SIGGRAPH 2004 Papers*, 2004, pp. 853–860.
- [16] T. Martin, E. Cohen, Volumetric parameterization of complex objects by respecting multiple materials, *Computers & Graphics* 34 (3) (2010) 187–197.
- [17] B. Li, X. Li, K. Wang, H. Qin, Generalized polycube trivariate splines, in: *Shape Modeling International Conf. (SMI)*, 2010, pp. 261–265.
- [18] J. Wang, S. Zhao, X. Tong, S. Lin, Z. Lin, Y. Dong, B. Guo, H.-Y. Shum, Modeling and rendering of heterogeneous translucent materials using the diffusion equation, *ACM Transactions on Graphics* 27 (1) (2008) 9:1–18.
- [19] X. Yin, M. Jin, F. Luo, X. D. Gu, Discrete curvature flows for surfaces and 3-manifolds, in: *Emerging Trends in Visual Computing*, Springer-Verlag, 2009, pp. 38–74.
- [20] J. Gregson, A. Sheffer, E. Zhang, All-hex mesh generation via volumetric polycube deformation, *Computer Graphics Forum* 30 (5) (2011) 1407–1416.
- [21] M. Nieser, U. Reitebuch, K. Polthier, CubeCover - Parameterization of 3D volumes, *Computer Graphics Forum* 30 (5) (2011) 1397–1406.
- [22] F. Kälberer, M. Nieser, K. Polthier, QuadCover - Surface parameterization using branched coverings, *Computer Graphics Forum* 26 (3) (2007) 375–384.
- [23] B. Lévy, Y. Liu, L_p centroidal voronoi tessellation and its applications, *ACM Transactions on Graphics* 29 (4) (2010) 119:1–11.
- [24] S. Han, J. Xia, Y. He, Hexahedral shell mesh construction via volumetric polycube map, in: *ACM Symposium on Solid and Physical Modeling*, SPM '10, 2010, pp. 127–136.
- [25] W. Zeng, J. Marino, A. Kaufman, X. D. Gu, Volumetric colon wall unfolding using harmonic differentials, *Computers & Graphics* 35 (3) (2011) 726–732, *Shape Modeling International (SMI) Conf.* 2011.
- [26] M.-F. Li, S.-H. Liao, R.-F. Tong, Facial hexahedral mesh transferring by volumetric mapping based on harmonic fields, *Computers & Graphics* 35 (1) (2011) 92–98.
- [27] Y. Lipman, D. Levin, D. Cohen-Or, Green coordinates, *ACM Transactions on Graphics* 27 (2008) 78:1–10.
- [28] M. Müller, J. Dorsey, L. McMillan, R. Jagnow, B. Cutler, Stable real-time deformations, in: *ACM Symposium on Computer Animation*, SCA '02, 2002, pp. 49–54.
- [29] S. Lang, *Algebra*, Graduate texts in mathematics, Springer, 2002.
- [30] B. Lévy, S. Petitjean, N. Ray, J. Maillot, Least squares conformal maps for automatic texture atlas generation, in: *ACM SIGGRAPH '02*, 2002, pp. 362–371.
- [31] Graphite software, <http://alice.loria.fr> (Feb. 2012).
- [32] J. Solomon, M. Ben-Chen, A. Butscher, L. Guibas, As-killing-as-possible vector fields for planar deformation, *Computer Graphics Forum* 30 (5) (2011) 1543–1552.
- [33] P. Joshi, M. Meyer, T. DeRose, B. Green, T. Sanocki, Harmonic coordinates for character articulation, in: *ACM SIGGRAPH'07*, ACM, 2007.
- [34] T. Ju, S. Schaefer, J. Warren, Mean value coordinates for closed triangular meshes, in: *ACM SIGGRAPH'05*, ACM, 2005, pp. 561–566.