

Reversible Harmonic Maps between Discrete Surfaces

DANIELLE EZUZ, Technion - Israel Institute of Technology

JUSTIN SOLOMON, Massachusetts Institute of Technology

MIRELA BEN-CHEN, Technion - Israel Institute of Technology

Information transfer between triangle meshes is of great importance in computer graphics and geometry processing. To facilitate this process, a *smooth and accurate map* is typically required between the two meshes. While such maps can sometimes be computed between nearly-isometric meshes, the more general case of meshes with diverse geometries remains challenging. We propose a novel approach for *direct* map computation between triangle meshes without mapping to an intermediate domain, which optimizes for the *harmonicity* and *reversibility* of the forward and backward maps. Our method is general both in the information it can receive as input, e.g. point landmarks, a dense map or a functional map, and in the diversity of the geometries to which it can be applied. We demonstrate that our maps exhibit lower conformal distortion than the state-of-the-art, while succeeding in correctly mapping key features of the input shapes.

CCS Concepts: • **Computing methodologies** → **Shape analysis**;

ACM Reference Format:

Danielle Ezuz, Justin Solomon, and Mirela Ben-Chen. 2019. Reversible Harmonic Maps between Discrete Surfaces. *ACM Trans. Graph.* 1, 1, Article 1 (January 2019), 13 pages. <https://doi.org/10.1145/3202660>

1 INTRODUCTION

Mapping 3D shapes to one another is a basic task in computer graphics and geometry processing. Correspondence is needed, for example, to transfer artist-generated assets such as texture and pose from one mesh to another [Sumner and Popović 2004], to compute in-between shapes using shape interpolation [Heeren et al. 2012; Von-Tycowicz et al. 2015], and to carry out statistical shape analysis [Munsell et al. 2008]. In these applications, desirable correspondences satisfy some basic key properties: They should be *smooth* to avoid introducing geometric noise during transfer; they should *preserve semantic features* to ensure that key features are put in correspondence; and they should be *reversible*, namely invariant to which of the two shapes is chosen as the source.

Many existing approaches to shape mapping focus on generating maps with low *global* distortion (e.g. preserving pairwise distances [Sahillioglu and Yemez 2011]) at the expense of large local distortion, which reduces the quality of the correspondence and hinders downstream applications. On the other hand, approaches that minimize local distortion measures mostly require an intermediate domain and construct the final map as a composition through this

domain (e.g. [Aigerman and Lipman 2016]). While such methods minimize distortion of the maps into the intermediate domain, the distortion of the composed map can be large. This problem is exacerbated when the input shapes have significantly different geometric features, such as four-legged animals with different dimensions, e.g. a cat and a giraffe. In this case, the *isometric distortion* of the optimal map is expected to be large, and thus minimizing the distortion of the two maps into an intermediate domain is quite different from minimizing the distortion of the composition.

We propose a novel approach for computing a smooth and reversible map between surfaces that are not isometric to each other, without requiring an intermediate domain. We incorporate semantic information by starting from some user guidance given in the form of sparse landmark constraints or a functional correspondence. Our main contribution is the formulation of an optimization problem whose objective is to minimize the *geodesic Dirichlet energy* of the forward and backward maps, while maximizing their reversibility. We compute an approximate solution to this problem using a high-dimensional Euclidean embedding and an optimization technique known as *half-quadratic splitting* [Geman and Yang 1995]. We demonstrate that our maps have considerably lower local distortion than those from state-of-the-art methods for the difficult case of non-isometric deformations. We further show that our maps are semantically accurate by measuring their adherence to self-symmetries of the input shapes, their agreement with ground-truth when the deformation is known, and their compatibility with human-generated segmentations.

1.1 Related Work

The shape correspondence literature is vast, and we refer the reader to recent surveys for a thorough review of the state-of-the-art in geometry-driven [Tam et al. 2013; Van Kaick et al. 2011] and data-driven [Xu et al. 2016] shape correspondence. We will focus our related work overview on methods for computing maps between triangular meshes that can handle shapes that are far from being related by an isometric deformation. In this realm, we characterize methods by the type of input they require and the type of output they generate. We therefore distinguish between *vertex-to-vertex* maps that yield a correspondence between the source and target vertices of the triangular meshes, *precise* maps that map every vertex on the source to a *point* on the triangulated surface of the target, and *generalized* maps that put in correspondence functions or probability distributions.

Fully-automatic methods. Kim et al. [2011] suggested one of the first fully automatic methods for non-isometric shape matching, that consistently generated high-quality outputs on a benchmark of shapes. This approach, denoted by BIM, generates a *precise* map as a blending of conformal maps, with blending weights optimized to

Authors' addresses: Danielle Ezuz, Technion - Israel Institute of Technology, danielle.ezuz@gmail.com; Justin Solomon, Massachusetts Institute of Technology, jsolomon@mit.edu; Mirela Ben-Chen, Technion - Israel Institute of Technology, mirela@cs.technion.ac.il.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2019 Copyright held by the owner/author(s).

0730-0301/2019/1-ART1

<https://doi.org/10.1145/3202660>

minimize isometric distortion. While providing excellent results in many cases, BIM is limited to genus-zero surfaces and can introduce large distortions for some shapes. Recently, Zheng et al. [2017] suggested to map between high-genus surfaces with the same genus, by decomposing the surfaces using a pants decomposition, and then computing harmonic maps between a set of intermediate cylindrical domains. This leads to a piecewise harmonic map between the input surfaces, which is further relaxed using geodesic heat flow. While this approach can be used without user intervention, if a globally semantic map is needed then accurate input landmarks are required, and it is limited to shapes of the same genus. Finally, Löhner et al. [2017] suggested a method that computes *vertex-to-vertex* maps based on a set of matching descriptors and pairwise distances. While the method is robust to topological changes, it may leave areas unmapped, and is therefore less appropriate for applications such as texture transfer.

The functional map approach, introduced by Ovsjanikov et al. [2012], was originally designed for nearly-isometric shapes but has since been extended to non-isometric matching [Burghard et al. 2017; Kovnatsky et al. 2015, 2013]. This method generates a generalized map which puts in correspondence the function spaces on the mapped shapes. Nogneng and Ovsjanikov [2017] recently suggested a method that computes functional maps using fewer input descriptors by formulating commutativity constraints, and Huang and Ovsjanikov [2017] suggested to use *adjoint* functional maps to improve and analyze correspondences. Using a hybrid approach, Maron et al. [2016] optimize jointly for a functional map and sparse correspondences, which are then extended to a dense vertex-to-vertex map. Finally, Solomon et al. [2016] compute a generalized map, which puts in correspondence probability distributions on the input shapes, by minimizing the Gromov–Wasserstein distance between the shapes.

While generalized maps are beneficial for challenging mapping problems, such as mapping between shapes of different genus, extracting a precise map is a necessary post-processing step if the output map is to be used for transferring high-frequency data such as textures, normals, or deformations. Furthermore, when the shapes are geometrically different and the optimal map is not expected to be isometric, the shape correspondence problem is ill-posed without additional *semantic* information. Such information can be given in the form of landmark constraints or an initial generalized map, from which a refined dense map can be computed.

Input: landmarks. Parameterization-based approaches compute bijective smooth maps to a common intermediate domain, and define precise maps between arbitrary shapes as the composition of the maps to the common domain. A variety of intermediate domains have been used in the literature, e.g. the plane [Aigerman et al. 2015], the sphere [Gu et al. 2004], the hyperbolic disk [Shi et al. 2017] and orbifolds [Aigerman et al. 2017; Aigerman and Lipman 2015, 2016; Tsui et al. 2013], to mention a few. These methods optimize the distortion of the map from the shapes to the target domain, but the composed map is not guaranteed in general to have low distortion. Furthermore, mapping through an intermediate domain places a topological restriction on the type of mapped shapes, as they should be topologically equivalent. Alternatively, Panozzo

et al. [2013] compute a direct map between two triangle meshes without requiring an intermediate domain. Their method computes on-surface barycentric coordinates with respect to the source landmarks, and then uses them with respect to the target landmarks to compute the corresponding point on the target shape. Similar to our approach, they use a high-dimensional Euclidean embedding to speed up the computation. Despite excellent results, in some cases a large number of landmarks is required to compute a correct correspondence. More recently, Mandad et al. [2017] use landmarks or extrinsic alignment for initialization and then optimize simultaneously for a generalized map and a precise map. We compare with their approach and demonstrate that our maps achieve better conformal distortion and adherence to the shape semantics.

Input: generalized maps. Several methods for recovery of vertex-to-vertex maps from generalized maps have been suggested. Shtern and Kimmel [2014] suggest a refinement technique that is based on heat kernel alignment, Rodolà et al. [2015] use a probabilistic model, and Vestner et al. [2017] solve a linear assignment problem. In general, vertex-to-vertex maps have higher conformal distortion than precise maps. Alternatively, Ezuz and Ben-Chen [2017] suggest a pointwise recovery method that generates *precise* maps using a smoothness prior, based on a spectral approach. Our method does not rely on spectral representations, and thus exhibits fewer artifacts in complex cases.

1.2 Contributions

We present an algorithm for shape correspondence between non-isometric triangular meshes, that has the following advantages:

- The algorithm is widely applicable, and the resulting maps are semantic and exhibit low conformal distortion.
- The formulation is simple and efficient to optimize and thus can be combined with additional energy terms and various initializations.
- The maps are accurate enough for downstream applications, such as shape interpolation and quad mesh transfer.

2 BACKGROUND: HARMONIC MAPS AND LOCAL DISTORTION

Suppose $M_1, M_2 \subseteq \mathbb{R}^3$ are smooth, compact surfaces with or without boundary. Given a map $\phi_{12} : M_1 \rightarrow M_2$ from one into another, a natural task is to measure the distortion of M_1 as it is mapped via ϕ_{12} onto M_2 ; this distortion measure eventually will serve as an objective function for optimization problems whose unknown is the correspondence ϕ_{12} . The basic role of these distortion measures is to evaluate whether nearby points are mapped to nearby points under ϕ_{12} , at least differentially, a common proxy for the quality of the map.

In the theory of differential geometry, a key distortion measure is the *Dirichlet energy* $E[\cdot]$ (defined below) of ϕ_{12} ; minimizers of $E[\cdot]$ are called *harmonic maps*. Intuitively, if we think of M_1 as a rubber sheet, a harmonic map represents an equilibrium position of the sheet after stretching it over M_2 and letting it compress. The Dirichlet energy and its minimizers find many roles in the geometry processing literature, most prominently in surface parameterization [Lévy et al. 2002], due to its intuitive measurement of distortion and connections

to notions of conformality. At the same time, theory and practice of harmonic mapping become considerably more challenging when M_2 has areas of positive curvature; intuitively these can cause the rubber sheet to slip or bunch, yielding singularities in gradient flow procedures designed to uncover harmonic maps.

In this section, we describe the basic construction of the Dirichlet energy and point out its advantages and flaws in the context of surface-to-surface correspondence; we also provide basic constructions for approximating the Dirichlet energy of a map between discrete surfaces. In §3, we then propose a modified notion of harmonicity designed to avoid singularities and asymmetry in the surface-to-surface correspondence pipeline.

2.1 Smooth Surfaces

Following [Nishikawa 2000; Urakawa 1993], harmonic maps between smooth surfaces are defined as the critical points of the *Dirichlet energy*:

$$E[\phi_{12}] := \frac{1}{2} \int_{M_1} |d\phi_{12}|^2 dv_1, \quad (1)$$

where $d\phi_{12}$ is the *map differential* and dv_1 is the volume element of M_1 . $E[\phi_{12}]$ measures the total stretch of M_1 after it is warped onto M_2 , as measured by the integrated norm of the Jacobian $d\phi_{12}$. Formally, given an orthonormal basis $\{e_1, e_2\}$ for $T_p M_1$ at $p \in M_1$, the integrand can be expanded as

$$|d\phi_{12}|^2 = \sum_{i=1}^2 \langle d\phi_{12}(e_i), d\phi_{12}(e_i) \rangle_{g_2(\phi_{12}(p))},$$

where g_2 is the metric of M_2 .

Existence, uniqueness, and regularity of harmonic maps given assumptions on the geometry/topology of M_1 and M_2 as well as the homotopy class of ϕ_{12} is a key theme in the twentieth-century differential geometry literature. A landmark paper by Eells and Sampson [1964] proves existence of a harmonic map in each homotopy class under the assumption that M_2 has non-positive curvature. The proof technique in this paper is attractive from a computational perspective: Essentially they start with an arbitrary map in the prescribed homotopy class and use an analog of gradient descent to decrease the Dirichlet energy.

A key drawback of the Eells and Sampson proof technique from a computational perspective, however, highlights an issue with harmonic correspondence in the context of algorithmic mapping between surfaces. In particular, their gradient descent procedure can fail when the target M_2 has regions of positive curvature. Roughly,




Fig. 1. Limitations of the piecewise linear discretization of the Dirichlet energy. (a) Source M_1 : a flat disk embedded in \mathbb{R}^3 . (b) Target M_2 : enneper. (c) Initial piecewise linear map. (d,e) Final maps that minimize the energies in Eqs. (3) and Eq. (4), respectively. See the text for details.

this singular behavior is explained by the fact that the objective $|d\phi_{12}|^2$ is minimized globally by $d\phi_{12} \equiv 0$, the constant map! This observation highlights the difference between harmonic mapping and elastic models like the ones proposed in [Chao et al. 2010; Sorkine and Alexa 2007], which seek $d\phi_{12}$ to be close to a rotation matrix rather than to the zero matrix. We will address this issue in our “reversible harmonic” formulation by adding the Dirichlet energy of ϕ_{12}^{-1} for the case of diffeomorphic correspondence; this has the added benefit of making forward maps ϕ_{12} and reverse maps ϕ_{21} critical points of the same objective function.

2.2 Triangle Meshes

For surfaces that are discretized as triangle meshes M_i , represented by their vertex, edge and face sets $(\mathcal{V}_i, \mathcal{E}_i, \mathcal{F}_i)$, a pointwise vertex map ϕ_{12} assigns a point on a face of M_2 to each vertex of M_1 . The extension of the vertex map to the interior of faces of M_1 determines the corresponding Dirichlet energy.

If the map is assumed to be *affine* on every face $f \in \mathcal{F}_1$, then the Dirichlet energy is given by

$$E[\phi_{12}] = \frac{1}{2} \sum_{f \in \mathcal{F}_1} |d\phi_{12}(f)|_2^2 a_f, \quad (2)$$

where $d\phi_{12}(f) \in \mathbb{R}^{2 \times 2}$ is the unique linear transformation between f and its image triangle $\phi_{12}(f)$, and a_f is the area of f [Pinkall and Polthier 1993]. Equivalently, the energy can also be written as:

$$E[\phi_{12}] := \frac{1}{4} \sum_{(u,v) \in \mathcal{E}_1} w_{uv} \|\phi_{12}(u) - \phi_{12}(v)\|_2^2, \quad (3)$$

where w_{uv} is the *cotangent weight* of the edge (u, v) . This energy is convex and quadratic in the images of the vertices of M_1 , and is therefore straightforward to minimize efficiently when ϕ_{12} is unrestricted, e.g. for planar parameterization [Lévy et al. 2002].

When M_2 is a non-Euclidean space, ϕ_{12} should be restricted to lie on M_2 , leading to a constrained optimization problem that is harder to solve. In addition, a more serious issue is the linearity assumption itself. When ϕ_{12} does not sample the target surface M_2 well, the linear extension of $\phi_{12}(\mathcal{V}_1)$ can be far from M_2 . In this case, minimizing the Dirichlet energy of the piecewise-affine map can lead to incorrect results.

Consider for example, as in Figure 1, mapping a disk M_1 (a) to an enneper surface M_2 (b) with Dirichlet boundary conditions; since the target has negative curvature, in the smooth case [Eells and Sampson 1964] gradient flow will reach a harmonic map $\phi_{12} : M_1 \rightarrow M_2$. An initial map (c) maps all the interior vertices of M_1 to a single interior vertex on M_2 , and the boundary of the disk is mapped to the boundary of the enneper. Minimizing Eq. (3) using gradient descent, the analog of Eells & Sampson’s heat flow, with the side constraint that $\phi_{12}(\mathcal{V}_1)$ is restricted to lie on M_2 , leads to a map (d) that is clearly not smooth. Effectively, Eq. (3) aims to place the image of every vertex of M_1 in the Euclidean weighted average of the image of its neighbors. When the affine map samples the target poorly, this strategy fails to generate an approximation of a smooth map.

Alternatively, Izeki and Nayatani [2005] suggest an intrinsic formulation, the *geodesic harmonic energy*, which replaces the Euclidean distances in Eq. (3) with the geodesic distances $d_{M_2}(\cdot, \cdot)$,

as follows:

$$E_D[\phi_{12}] := \sum_{(u,v) \in \mathcal{E}_1} w_{uv} d_{M_2}^2(\phi_{12}(u), \phi_{12}(v)). \quad (4)$$

As shown in Figure 1(e), minimizing this energy instead of the Euclidean one yields a significantly better result at the cost of having to compute geodesic distances. Motivated by this idea, we propose to use this energy as the main building block in a shape mapping algorithm. We reformulate it to allow efficient optimization and combination with other terms that address the case of positively-curved target surfaces M_2 , as described in the following section.

3 REVERSIBLE HARMONIC MAPS

Notation. We represent a triangle mesh M by its vertex, edge and face sets $(\mathcal{V}, \mathcal{E}, \mathcal{F})$, respectively, where we denote $n = |\mathcal{V}|$, and its given embedding by $V \in \mathbb{R}^{n \times 3}$. We denote scalar functions $g: M \rightarrow \mathbb{R}$ by a vector of coefficients of piecewise linear hat functions, with $g \in \mathbb{R}^n$. The squared l_2 norm of a function on the surface is given by $\|g\|_M^2 = g^T A g$, where $A \in \mathbb{R}^{n \times n}$ is the diagonal (lumped) mass matrix of the vertices. The total area of the mesh is denoted by $s = \text{tr}(A)$. The squared gradient norm is given by $\|g\|_W^2 = g^T W g$, where W is the matrix of cotangent weights. Similarly, for matrices $G \in \mathbb{R}^{n \times k}$ whose columns are scalar functions, we use the matrix trace: $\|G\|_M^2 = \text{tr}(G^T A G)$, $\|G\|_W^2 = \text{tr}(G^T W G)$. When two meshes are involved we use a subscript, e.g. A_i is the mass matrix of M_i .

3.1 Energy

While the geodesic harmonic energy can be effective, harmonic maps in general can become degenerate and map large regions to a single point. Indeed, the map taking all the points on M_1 to a single point on M_2 is harmonic. An example of such behavior is demonstrated in Figure 2. On the top row, the source M_1 is a sphere with a small number of triangles (a), that is mapped to a

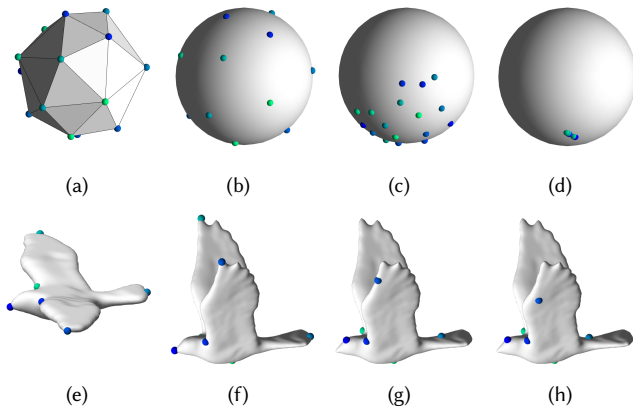


Fig. 2. Collapse of a harmonic map. Top row: mapping from a low resolution sphere (a) to a high resolution sphere, starting from the identity map (b). The map quickly “slides” to a single hemisphere (c) and then degenerates (d). Bottom row: the same phenomenon with more complex shapes from SHREC’07 [Giorgi et al. 2007], where we use a sparse set of landmarks for initialization and visualization. The final result (h) does not map any points to the upper part of the wings and to most of the tail.

high-resolution target sphere M_2 . Even if the initial map is the ground-truth map (b) between the spheres, during the optimization the map quickly “slides over” to a single hemisphere of the target sphere (c), and then degenerates and collapses to a single point (d). The same phenomenon occurs for complex shapes, as shown in the bottom row. Here, we initialized a map between two bird shapes from SHREC’07 [Giorgi et al. 2007] (e,f) using landmarks as specified in section 4.2. Initially (f), the tips of the wings and tail are mapped correctly, but again they gradually slide (g) until no vertices are mapped to most of the target’s wings or tail (h).

Intuitively, in the discrete case, we can think of M_1 as an elastic fishnet instead of a continuous rubber sheet, stretched over M_2 and allowed to compress. Then, in addition to the usual degeneracies in the smooth case, the target surface can effectively “slip through” one of the holes in the net, allowing the map to degenerate to a single point. To prevent this, we minimize the geodesic harmonic energies of both the *forward* and *backward* maps, together with a *reversibility* constraint relating both maps. As we later demonstrate, this approach is highly effective in generating non-degenerate harmonic maps.

Smoothness. Given two triangle meshes M_1, M_2 , and maps ϕ_{12}, ϕ_{21} , the total harmonic energy of both forward and backward maps is given by

$$E_D[\phi_{12}, \phi_{21}] = \sum_{\substack{i,j \in \{1,2\} \\ i \neq j}} \frac{1}{s_j} E_D[\phi_{ij}], \quad (5)$$

where E_D is given in Equation (4).

Reversibility. We define the reversibility energy similarly to [Ezuz and Ben-Chen 2017; Kovnatsky et al. 2013] as:

$$E_R[\phi_{12}, \phi_{21}] = \sum_{\substack{i,j \in \{1,2\} \\ i \neq j}} \frac{1}{s_i} \sum_{p_i \in \mathcal{V}_i} d_{M_i}^2(\phi_{ji}(\phi_{ij}(p_i)), p_i) A_i(p_i). \quad (6)$$

The reversibility energy prevents the maps from collapsing. In the smooth case, if the reversibility energy is bounded pointwise, it easily follows that the maps are close to being injective and surjective, as we show in the Appendix. For both energies, care is required to handle correctly meshes of different scales, hence the normalization by s_i , the total area of M_i .

Finally, the full energy is given by:

$$E[\phi_{12}, \phi_{21}] = \alpha E_D[\phi_{12}, \phi_{21}] + (1 - \alpha) E_R[\phi_{12}, \phi_{21}], \quad (7)$$

where the parameter $\alpha \in [0, 1]$ controls the trade-off between smoothness and reversibility. Note that while in the continuous setting exact reversibility might be desired, this will not be the case in the discrete setting described in section 3.2. Thus we enforce reversibility as a soft constraint controlled by the parameter α .

To demonstrate the effect of the different components of the energy we compute a map between a disk and a folded disk, using different values of α . In the initial map all the interior vertices of the disk are mapped to a single interior vertex of the folded disk, and the boundary of the source is mapped to the boundary of the target. During the optimization the mapped boundary vertices are not constrained to lie on the boundary of the target.




Fig. 3. Preventing collapse with reversibility. For different α values, we measure the discrete geodesic Dirichlet energy and the sum of relative mapped area (ideally 2). We visualize some of the results using texture transfer (left), and show the final values as a function of α (right). Note that when α is small the Dirichlet energy is high, and when α is large the map collapses, as is evident by the zero total area. Finally, taking $\alpha = 5 \cdot 10^{-4}$ leads to a good balance between the energy components.

Figure 3 shows the results, where we visualize the maps using texture transfer (left), and quantitatively evaluate them using the discrete geodesic Dirichlet energy from Eq. (4) (center right) as well as the sum of the total area of the images of the forward and backward maps (right). The graphs show the final values of these quantities as a function of the parameter α , where the energy is normalized with respect to its value in the first iteration.

The figure demonstrates the trade-off that the parameter α controls, e.g. taking $\alpha = 0$ models reversibility only, leading to a high Dirichlet energy. On the other hand, taking $\alpha = 1$ models harmonicity only, and leads to a map that collapses the image to a single point, as is evidenced by the final total area which is zero. In general, if α is too small, the map has a high Dirichlet energy and thus is more distorted locally, and if α is too large the map collapses. Taking $\alpha = 5 \cdot 10^{-4}$, as we did for all the experiments except for Figure 3, leads to a balance between the harmonic energy and reversibility.

Minimizing the energy in Eq. (7) requires computing the gradient of the geodesic distances with respect to the forward and backward maps, as well as tracing vector fields on the surface, which are both computationally heavy. We therefore apply two approximations to address these issues.

3.2 Energy approximation

3.2.1 Notation. Any point $p \in M$ can be represented uniquely using its barycentric coordinates $\omega_l(p)$, $l \in \{1, \dots, 3\}$ with respect to the face $f(p) = (v_1, v_2, v_3) \in \mathcal{F}$ it lies on. We denote by $\lambda(p) \in \mathbb{R}^{1 \times n}$ the row vector that is zero everywhere except at the vertices of $f(p)$, where we have $\lambda(p)[v_l] = \omega_l(p)$. In addition, we denote the *feasible row set* of M , i.e. the set of all possible such vectors, by $\mathcal{P} = \{\lambda(p) \mid p \in M\}$. Finally, the *feasible set* of all possible precise maps from M_1 to M_2 is given by $\mathcal{P}_{12} = \{P_{12} \in \mathbb{R}^{n_1 \times n_2} \mid P_{12}(l, :) \in \mathcal{P}_2, \forall l \in \{1..n_1\}\}$, where $P(l, :)$ denotes the l -th row of the matrix P . Thus, any map ϕ_{12} can be represented using a matrix P_{12} , by setting $P_{12}(l, :) = \lambda_2(\phi_{12}(v_l))$, $\forall l \in \{1..n_1\}$, which, by definition, is in the feasible set \mathcal{P}_{12} . Furthermore, the matrix $P_{12}V_2 \in \mathbb{R}^{n_1 \times 3}$ represents the images of the vertices \mathcal{V}_1 under the map ϕ_{12} .

3.2.2 High-dimensional embedding. As we have seen, if the target space is *Euclidean* then the geodesic distances are Euclidean distances, and the optimization is simple and efficient. Following similar ideas in the literature [Bronstein et al. 2006], we therefore

suggest to use a *high-dimensional Euclidean embedding* as a proxy for fast geodesic distance computation.

Given a mesh M , we seek an embedding $x : \mathcal{V} \rightarrow \mathbb{R}^m$, for $m \ll n$ such that the Euclidean distance $\|x(u) - x(v)\|$ approximates well the geodesic distance $d_M(u, v)$, for all $u, v \in \mathcal{V}$. The literature on such embeddings is quite vast, and we chose to leverage the method suggested by Panozzo et al. [2013] that relies on multidimensional scaling [Cox and Cox 2000]. Any other embedding method could be used as well, as long as the geodesic distances are well approximated. We took $m = 8$ in all our experiments, following Panozzo et al. [2013].

Our goal now is to compute a harmonic map between the high-dimensional Euclidean embeddings. We denote by $X \in \mathbb{R}^{n \times m}$ the matrix whose rows are the embeddings $x(v)$, $\forall v \in \mathcal{V}$. Rewriting the harmonic and reversibility energies in terms of the high-dimensional embeddings, and in matrix form, leads to:

$$E(P_{12}, P_{21}) = \sum_{\substack{i, j \in \{1, 2\} \\ i \neq j}} \alpha \frac{1}{s_j} \|P_{ij} X_j\|_{W_i}^2 + (1 - \alpha) \frac{1}{s_i^2} \|P_{ij} P_{ji} X_i - X_i\|_{M_i}^2. \quad (8)$$

Note that the weights in the harmonic energy, given now in matrix form in W_i , remain the same for the high-dimensional embedding, since the embedding is nearly isometric.

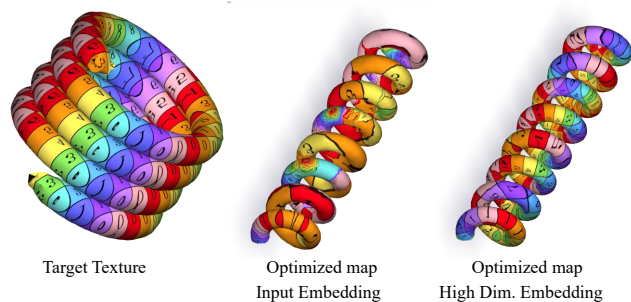


Fig. 4. The importance of the high dimensional embedding. For these shapes, geodesic and Euclidean distances are significantly different, thus using the input vertex positions in \mathbb{R}^3 during the optimization results in a highly distorted map (center). The high dimensional embedding that approximates the geodesic distances leads to a better map (right).

Figure 4 demonstrates the importance of the high dimensional embedding. For shapes where the geodesic distances are considerably different than Euclidean distances in \mathbb{R}^3 , e.g., the spring shapes from SHREC'07 [Giorgi et al. 2007], using the three-dimensional input Euclidean embedding in the optimization leads to a highly distorted map (center). Specifically, neighboring vertices on the source shape are mapped to different coils of the target spring, which are extrinsically close but intrinsically far. On the other hand, by using the high dimensional embedding, the geodesic Dirichlet energy is well approximated leading to an improved map (right).

3.2.3 Half quadratic splitting. While Equation (8) can be minimized using gradient descent, we found, similarly to [Ezuz and Ben-Chen 2017], that it is more efficient to use the *half quadratic splitting* optimization method [Geman and Yang 1995] (see also e.g. [Wang et al. 2008; Zoran and Weiss 2011]). We introduce auxiliary variables $X_{ij} \in \mathbb{R}^{n_i \times m}$, which estimate the images of the vertices \mathcal{V}_i given by $P_{ij}X_j$. Substituting, the energies are:

$$\bar{E}_D(X_{ij}) = \frac{1}{s_j} \|X_{ij}\|_{W_i}^2, \quad \bar{E}_R(P_{ij}, X_{ji}) = \frac{1}{s_i} \|P_{ij}X_{ji} - X_i\|_{M_i}^2, \quad (9)$$

for $i, j \in \{1, 2\}$, $i \neq j$. In addition, we need soft constraints for the auxiliary variables:

$$\bar{E}_Q(P_{ij}, X_{ij}) = \frac{1}{s_i s_j} \|X_{ij} - P_{ij}X_j\|_{M_i}^2, \quad (10)$$

where we again normalize by $s_i s_j$ to retain scale invariance.

The full energy is now:

$$\bar{E}(P_{12}, P_{21}, X_{12}, X_{21}) = \sum_{\substack{i, j \in \{1, 2\} \\ i \neq j}} \alpha \bar{E}_D(X_{ij}) + (1 - \alpha) \bar{E}_R(P_{ij}, X_{ji}) + \beta \bar{E}_Q(P_{ij}, X_{ij}), \quad (11)$$

where β controls the accuracy of the auxiliary variables and functions as a step size. When using the half-quadratic splitting optimization scheme, the update schedule for β is often tailored per application, with the general guideline of increasing β as the iterations advance [Wang et al. 2008]. In our case, we often initialize the optimization with a highly degenerate map, e.g. as obtained from a sparse set of landmarks, and therefore the value of β during the first iterations should be small enough so that the map can change significantly. As the iterative solution approaches a local optimum, β can increase, as less modification is required. The final value of β should be large enough to ensure P_{ij} and X_{ij} correspond. In all of our experiments, we took $\beta = 5 \cdot 10^{-3}k$ where k is the optimization iteration number for the first 100 iterations, and then kept the value of β fixed until convergence.

3.3 The optimization problem

Our optimization problem is now given by:

$$\begin{aligned} & \text{minimize}_{P_{12}, P_{21}, X_{12}, X_{21}} && \bar{E}(P_{12}, P_{21}, X_{12}, X_{21}) \\ & \text{subject to} && P_{12} \in \mathcal{P}_{12}, \quad P_{21} \in \mathcal{P}_{21}, \end{aligned} \quad (12)$$

where \mathcal{P}_{ij} is the feasible set of precise maps from M_i to M_j . Despite the two approximations that we used, solving this optimization problem succeeds in decreasing the total discrete Dirichlet energy

from Eq. (4) while preventing the map from collapsing, as is illustrated in Figure 5. In addition to the energy, we show the initial map that was created from landmarks as described in section 4.2, and the intermediate map at a few iterations.

4 OPTIMIZATION

Our optimization problem has block structure, in the sense that if some of the variables are kept fixed it becomes a linear least squares problem. We therefore chose to use block coordinate descent (see e.g. [Xu and Yin 2013]) as the optimization algorithm.

4.1 Block coordinate descent

In each sub-iteration we solve for one of the matrices X_{ij} , P_{ij} , while keeping the others fixed. Since the energy is quadratic in all the variables, every sub-iteration involves a relatively simple optimization problem, with the only complication arising because of the non-convex feasible sets \mathcal{P}_{ij} .

Optimizing for X_{ij} . When P_{12}, P_{21} are fixed, the optimization problem is a linear least squares minimization of the form $\|AX_{ij} - B\|_2^2$, with known matrices A and B , where A is sparse, which we solve using a direct method. The system is highly over-constrained, as even if P_{ji} degenerates, the system is well-conditioned due to the term \bar{E}_Q , as long as the vertex areas of the mesh M_i do not vanish.

Optimizing for P_{ij} . When X_{12}, X_{21} are fixed, the energy has the form $\|P_{ij}A - B\|_2^2$, where A, B are known, with the constraints that $P_{ij} \in \mathcal{P}_{ij}$. Following [Ezuz and Ben-Chen 2017], the optimization is done by solving for every row of P_{ij} separately. Intuitively, we can think of B as a high-dimensional embedding of the faces of M_j , and of A as a high-dimensional point cloud. The optimal P_{ij} projects each point in A to its closest point on the faces given by B . As shown in [Ezuz and Ben-Chen 2017], this process is guaranteed to find P_{ij} which are globally optimal when X_{ij} are kept fixed.

Stopping criterion. The alternating descent guarantees that the energy is reduced at every iteration, since the sub-iterations find a global optimum of the reduced optimization problems. In practice, we stopped the optimization when the change of energy was less than 10^{-9} , or after a maximum of $N = 200$ iterations. In most cases, we have observed convergence of the energy to high precision even when early stopping after N iterations was used.

We provide the details of the alternating descent in Algorithm 1.

ALGORITHM 1: Alternating minimization.

Input: Two triangles meshes M_1, M_2 , initial $P_{12}, P_{21}, X_{12}, X_{21}$ Output: P_{12}, P_{21} For $k = 1 \dots N$ For $i = 1, 2$ $j = 3 - i$ $P_{ij} \leftarrow \operatorname{argmin}_{P \in \mathcal{P}_{ij}} \bar{E}_R(P, X_{ji}) + \bar{E}_Q(P, X_{ij})$ $X_{ij} \leftarrow \operatorname{argmin}_{X \in \mathbb{R}^{n_i \times m}} \bar{E}_D(X) + \bar{E}_R(P_{ji}, X) + \bar{E}_Q(P_{ij}, X)$ end end

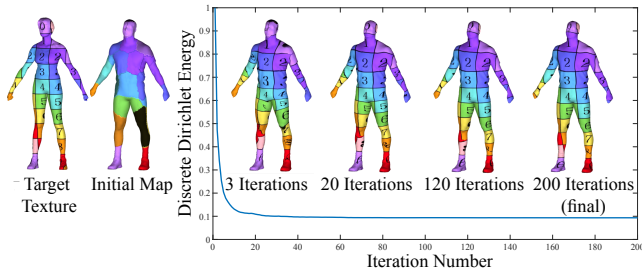


Fig. 5. Optimization of Eq. (12), starting from landmarks. The discrete Dirichlet energy in Eq. (4) decreases, and the final map, visualized by texture transfer, is semantic and not distorted locally.

4.2 Initialization

Our method is general and can receive as input various initial data. Depending on the input, we describe the initialization of the variables P_{ij}, X_{ij} . Given a pointwise map $\phi_{12} : \mathcal{V}_1 \rightarrow M_2$ its corresponding matrix representation is given by $P_{12}(l, :) = \lambda_2(\phi_{12}(v_l))$ for $l \in \{1, \dots, n_1\}$. Similarly, given a matrix representation $P_{ij} \in \mathcal{P}_{ij}$, we have that $\phi_{ij}(v_l) = (P_{ij}V_j)(l, :) \in M_j$. Therefore, in the following refer to P_{ij} or ϕ_{ij} according to which notation is more convenient.

Pointwise map. Given a pointwise map P_{12} we approximate an inverse map P_{21} by taking $\phi_{21}(v_2) = \operatorname{argmin}_{v \in \mathcal{V}_1} \|\phi_{12}(v) - v_2\|$. Then, the initialization of X_{ij} is $P_{ij}X_j$.

Functional map. The term *functional map* [Ovsjanikov et al. 2012] denotes a map between scalar functions. It is a linear operator that can be represented using a matrix when scalar functions are represented in a linear basis. Let $\Psi_i \in \mathbb{R}^{n_i \times k_i}$ be a matrix whose columns are basis functions of a subspace of scalar functions on M_i , where each function is piecewise linear and is defined by values assigned to vertices. Given a pointwise map that maps vertices of M_1 to points on M_2 , the corresponding functional map $C_{12} \in \mathbb{R}^{k_1 \times k_2}$ maps functions on M_2 to functions on M_1 , represented in the reduced basis. Therefore, given two functional maps C_{12} and C_{21} , we initialize $X_{ij} = \Psi_i C_{ij} \Psi_j^\dagger X_j$. Optimizing P_{ij} does not require initialization. Figure 9 shows results where functional maps were used for initialization.


Landmarks. Given r input landmark pairs $\{(p_i, q_i)\}$ where $p_i \in \mathcal{V}_1, q_i \in \mathcal{V}_2$ and $i \in \{1, \dots, r\}$, we first construct a rough initial pointwise map P_{12} and then use it to initialize the rest of the variables, as previously described. We first compute the geodesic Voronoi diagram on M_1 with centers p_i , and then set $\phi_{12}(v) = q_i, \forall v \in C_i$, where C_i is the geodesic cell corresponding to the center p_i . Note that this initialization is highly degenerate, as all the points on M_1 are mapped to the landmarks q_i on M_2 , yet it is enough for our needs. In Figure 5 initialization using input landmarks is visualized.

4.3 Limitations

Using the projection step when optimizing for P_{ij} has some limitations. First, as discussed in [Panozzo et al. 2013], such a projection is not smooth. Furthermore, the closest point on the high-dimensional embedding of the triangle mesh might not be unique, therefore the

solution of the optimization for P_{ij} might alternate between two configurations with the same energy. Thus, while the energy is guaranteed to converge, we do not have a similar guarantee for the convergence of the solution. In practice, we have not encountered a case where these limitations posed a practical problem. In future work it could be possible to handle the first issue using a Phong projection, as in [Panozzo et al. 2013], and the second issue using an additional regularization that penalizes diverting from the current solution.

Since we optimize for both harmonicity and reversibility we cannot guarantee convergence to a smooth harmonic map; this is likely a fruitful avenue for future work. In addition, our method does not consider extrinsic features such as edges and corners, and hence such features will not necessarily be mapped to each other. An example is shown in the inlined figure, where our map between the two folded disks is smooth, but the edge features do not correspond.



4.4 Timing

The most expensive step in the optimization process is the projection on a triangle mesh for optimizing P_{ij} . However, this procedure is highly parallelizable since the projection of each point is independent of the other points. We used CUDA 8 to implement the projection in parallel, while the rest of the optimization method was written in MATLAB. On a desktop machine with a TITANX GPU and an Intel Core i7 processor, 200 optimization iterations of our method, for shapes with 5K vertices, took around 115 seconds.

5 RESULTS

To validate our method we have compared with a variety of state-of-the-art mapping techniques, in accordance with the type of input they can accept. In addition, we show applications to shape interpolation and quad mesh transfer.

5.1 Quality metrics

To evaluate the quality of a map we measure its smoothness through its conformal distortion and its semantic accuracy, using the distance to the ground truth, when given. We also use alternative measures, such as symmetry and compatibility with ground truth segmentations, when no dense ground truth map is available.

Conformal distortion. We use the definition by Hormann and Greiner [2000, Eq. (3)] for the conformal distortion of a single triangle $f \in \mathcal{F}_1$: $\kappa(f) = \frac{\sigma_1}{\sigma_2} + \frac{\sigma_2}{\sigma_1}$, where $\sigma_1 \geq \sigma_2$ are the singular values of the linear transformation which maps f from M_1 to M_2 . We subtract 2 so that the minimal conformal distortion is zero and visualize the result as a cumulative graph showing the percentage of triangles with less than a certain distortion value.

Distance from ground truth. When a ground truth map is given, we measure the distance from the ground truth using the protocol suggested by Kim et al. [2011, Section 8.2]. For every mapped vertex, we measure its geodesic distance from the ground truth location,




Fig. 6. Quantitative comparison on the SHREC dataset, measuring, from left to right: conformal distortion, compatibility with symmetries and distance from ground truth landmarks. Note that we achieve a better conformal distortion, and comparable symmetry geodesic error. Furthermore, note that WA and HOT do not modify their input landmarks, while our method and VMTP do. Compared to VMTP we achieve a better landmark geodesic error.

relative to the square root of the total area of M_2 , and visualize the percent of vertices whose distortion is less than a given value.

Compatibility with segmentations. For some datasets ground truth labeled segmentations are available. In this case, for every pair of shapes and a given map we measure the consistency of the segmentation with respect to the map. This is done by computing the relative vertex area of vertices that are mapped to a face that belongs to the same segment as the source vertex.

Compatibility with symmetry. For some datasets a ground truth map is only known for a subset of the points, yet a full intrinsic symmetry can be computed for every shape separately. We assume that a good map should respect the intrinsic symmetries of the source and target shapes, given by S_1, S_2 , respectively. We therefore




Fig. 7. Qualitative results, from input landmarks. From left to right: target texture, our method, HOT [Aigerman and Lipman 2016], WA [Panozzo et al. 2013] and VMTP [Mandad et al. 2017]. See the text for details.

use these symmetries as input, and measure the compatibility of the map ϕ_{12} with the symmetries, given by the geodesic distance $d_{M_2}(S_2(\phi_{12}(v_1)), \phi_{12}(S_1(v_1)))$, $\forall v_1 \in \mathcal{V}_1$. We visualize the result using a cumulative graph, similarly to the ground truth error. To compute the symmetries we use the method by Kim et al. [2011]. We also manually filtered the results to use only the accurate symmetries.

5.2 Dataset: SHREC, input: landmarks

We use the BIM benchmark [Kim et al. 2011] that provides more than 200 pairs of highly non-isometric shapes from the SHREC dataset [Giorgi et al. 2007] with user-verified landmarks. We compare our method with a state of the art parameterization based method [Aigerman and Lipman 2016] (HOT) and the weighted averages method [Panozzo et al. 2013] (WA). Both receive as input landmark points, which are not modified during the optimization, and generate precise maps. In addition we compare to the recent method by Mandad et al. [2017] (VMTP), that similarly gets as input landmark points, yet can modify them during the optimization. Since VMTP requires uniform isotropic meshes, we recursively add edges using the longest edge bisection method to meshes with less than 10K vertices, before applying VMTP. All the methods we compare to produce *precise* maps, as vertex-to-vertex maps induce high local distortion. As input to our method we also use the user defined landmarks, and extend them to a full initial map as described in section 4.2. The landmarks are not used after the initialization.

Quantitative results are shown in Figure 6, where we measure conformal distortion, compatibility with symmetries and distance from the ground truth landmarks. Note that our method achieves the best conformal distortion. In addition, the distance from the ground truth landmarks is also improved when compared to the other method which modifies them (VMTP). As shown in section 5.3, the option to modify the input landmarks is valuable when the input is not completely reliable. In terms of compatibility with symmetry, our method is comparable with existing techniques, notably achieving a better ratio of perfect matches with about 15% of the vertices exactly symmetric for our method, where the next best method has less than 10% exactly symmetric vertices. On this dataset ground truth segmentations are also available [Chen et al. 2009; Kalogerakis et al. 2010], and measuring the relative mapped area which is compatible with the segmentations we have HOT: 90.39%, WA:

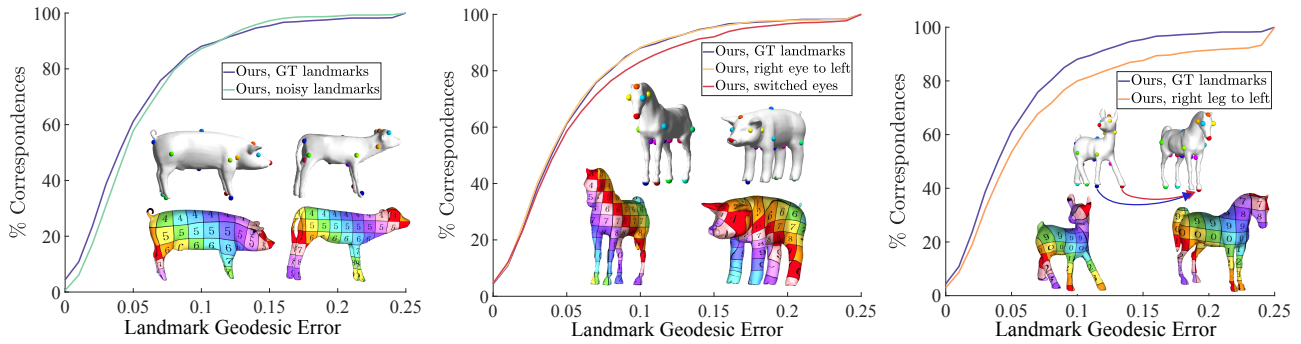


Fig. 8. Mapping SHREC quadrupeds by our method, starting from noisy landmarks. Our method is only slightly affected by the noise even when the landmark modification is severe.

90.35%, VMTP: 81.8%, our method: 89.62%. Therefore, this measure also demonstrates that our maps are as compatible semantically as existing techniques, while being considerably more conformal.

Qualitative results are shown in Figure 7, where we have selected a subset of pairs to visually show the differences between the maps. In every row we show, from left to right, the target texture, and the results of our method, HOT, WA and VMTP.

5.3 Dataset: SHREC quadrupeds, input: noisy landmarks

In many cases, the selection of the landmarks by the user has some variability (see, e.g. [Chen et al. 2012]), and it might be better to treat

these landmarks as guidelines rather than exact ground truth. Our approach is compatible with this notion, since our method only uses the landmarks for initialization, and their final location will, in most cases, vary from their initial one. To check the sensitivity of our approach to the landmark locations, we repeated the experiment from section 5.2 with various landmark modifications. Figure 8 shows the landmark geodesic error of the output maps when starting from the noisy landmarks compared to starting from the original landmarks. We ran the experiment on the “quadrupeds” class (20 pairs) from the SHREC dataset and did the following modifications: (a) moved every landmark randomly to a vertex in its 5-ring neighborhood, (b) switched between the two eyes, or mapped both eyes on M_1 to a single eye on M_2 , and (c) mapped both feet on M_1 to the same foot on M_2 . In addition to the error graph we show some example maps, as well as the input noisy landmarks. As the figure shows, our results are not sensitive to landmark noise, and even a relatively severe modification, such as mapping both feet to the same foot, yields good qualitative and quantitative results.

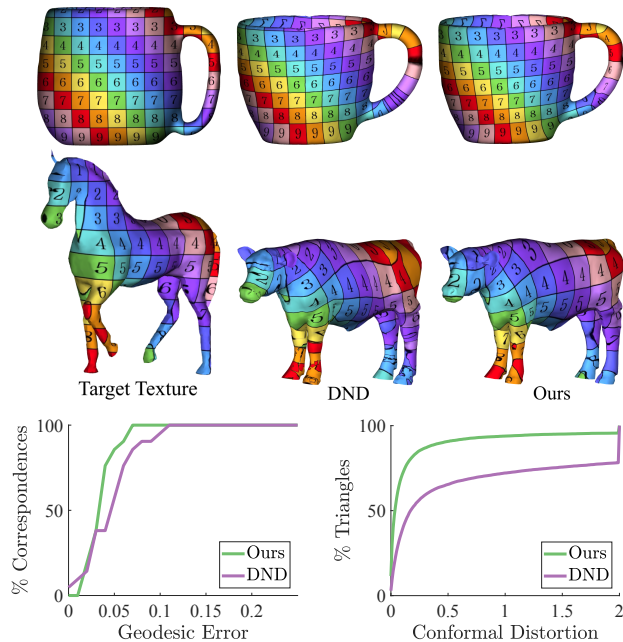


Fig. 9. Qualitative and quantitative comparison starting with a functional map computed from landmarks. From left to right: target texture, [Ezuz and Ben-Chen 2017] (DND), our method. Notice the difference at the cup handle and the legs.

5.4 Dataset: SHREC two pairs, input: functional map

The functional map [Ovsjanikov et al. 2012] machinery is quite versatile, and allows to compute generalized maps in a variety of cases. Our method can also be used to extract a precise pointwise map from a given functional map. We use the SHREC dataset with its landmark data from the BIM benchmark [Kim et al. 2011], and use the landmarks to compute a functional map using the Wave Kernel Map and the Wave Kernel Signature [Aubry et al. 2011]. Any other recent method for computing functional maps could be used as well. We provide the functional map as input to our approach and the recent map deblurring approach [Ezuz and Ben-Chen 2017] (DND), which is the only other method that recovers *precise* maps from a functional map. Specifically, we used the consistency extension of DND with $\alpha = 0.8$. Figure 9 qualitatively visualizes the difference between the methods for two pairs of shapes. Note the map improvement on the handle of the cup and the legs of the cow. We also show graphs of the conformal distortion and ground-truth error of the landmarks.

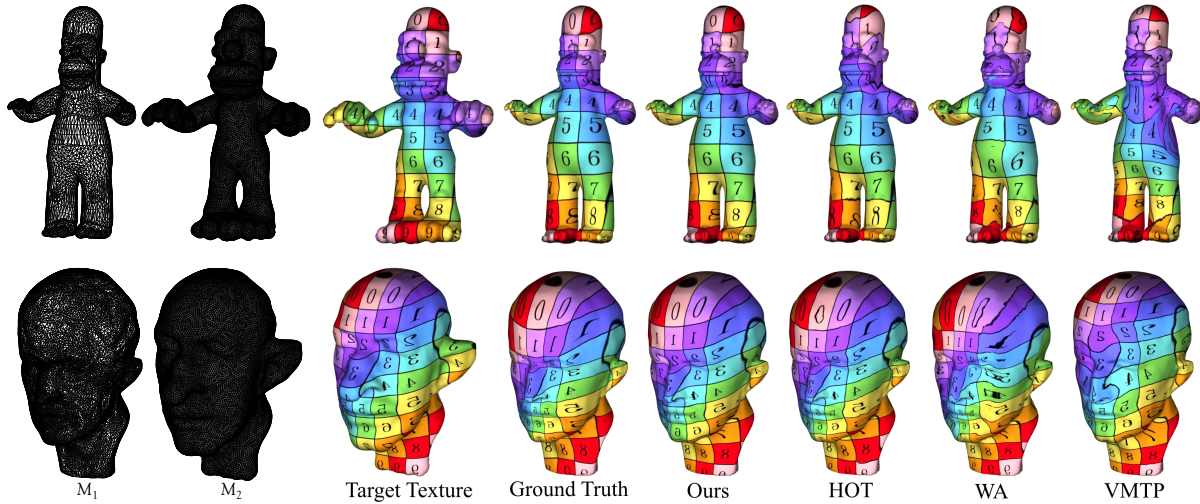


Fig. 10. Qualitative result, caricatures. M_2 was generated by deforming M_1 [Sela et al. 2015], and the deformation defines a ground truth map. M_2 was remeshed so that the source and target shapes do not have the same connectivity.

5.5 Dataset: caricatures, input: landmarks

One of the advantages of our formulation is its simplicity, that leads to flexibility in adding additional components to the energy. For example, in some cases it can be beneficial to add *weak* landmark constraints, to encourage feature points to remain in the neighborhood of the input landmarks.

Weak landmark constraints. Given pairs of matching landmarks $p_i \in \mathcal{V}_1, q_i \in \mathcal{V}_2, i = 1 \dots r$, we add the following term to the energy:

$$\gamma \sum_{i=1}^r A_1(p_i) \|X_{12}(p_i, \cdot) - X_2(q_i)\|_{M_2}^2 + A_2(q_i) \|X_{21}(q_i, \cdot) - X_1(p_i)\|_{M_1}^2. \quad (13)$$

The value of γ depends on the reliability of the landmarks, if the landmarks are not accurate then γ should be small. To demonstrate the effectiveness of this approach, we used as input the Homer and Max Planck models, and caricatures of these models generated using the method by Sela et al. [2015]. The caricatures have the same triangulation as the original shapes. We remesh the caricatures to avoid bias, and project the original caricature to the remeshed model to generate the ground truth for validation. As input to all




Fig. 11. Quantitative measures for the meshes in Figure 10. Note that our method achieves a considerably better conformal distortion, while maintaining ground truth error comparable to existing methods.

approaches, we picked 28 and 14 landmarks from the input map for the homer and Max models respectively. For our algorithm we added the energy in Equation (13) with $\gamma = 1$. The qualitative and quantitative results are shown in figures 10 and 11, respectively. As in previous experiments, we achieve considerably better conformal distortion, with comparable ground truth error.

5.6 Robustness

We test our method on horse shapes from the SHREC'10 dataset [Bronstein et al. 2010], which contains various shapes with geometric and topological noise. The results are shown in the top row of Figure 12,

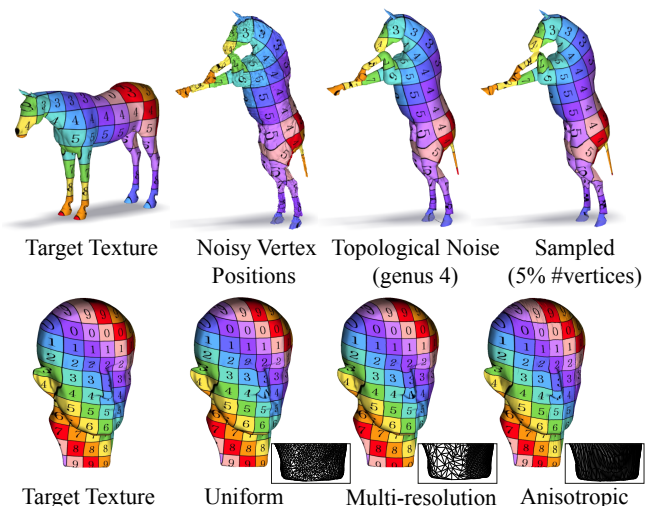


Fig. 12. Robustness to noise and sampling. Top row: a shape with various transformations from SHREC'10 [Bronstein et al. 2010]. Bottom row: the left shape is mapped to the same geometry with different tessellations.




Fig. 13. Shape interpolation using our computed correspondence as input for [Heeren et al. 2014].

where we compute a correspondence between the horse in a neutral pose (left) and a noisy pose: (center left) noise is added to the vertex positions; (center right) topological noise is introduced, e.g. between the rear feet; and (right) the mesh is sampled to 5% of the number of vertices. We initialized our method using 19 landmarks and obtained semantically correct results in all cases. The bottom row shows our results for different tessellations of a shape of a head, where one shape has uniform triangles, a second shape has larger triangles on one half and smaller triangles on the other half, and a third shape has highly anisotropic triangles. For initialization, we used the identity map with additional noise. The resulting mapped texture is smooth and similar for the different tessellations.

5.7 Application: Shape Interpolation

Existing shape interpolation methods require the input shapes to share the same connectivity, while real data rarely satisfies this requirement. Our mapping can be used to *remesh* the target surface M_2 using the image of the vertices of M_1 , given by $P_{12}V_2$, and the connectivity of M_1 . We used our map between two birds from SHREC to demonstrate this application, starting from the BIM landmarks as described in section 5.2. After remeshing M_2 , as a few faces had a zero area, we iteratively moved vertices of the degenerate faces to an average of their 1-ring neighbors until there were no degenerate faces. We then used the shape interpolation method by Heeren et al. [2014] to interpolate between the shapes, as shown in Figure 13. Note that we correctly mapped the wings, head and tail of the birds, as is evident from the natural interpolation results.

5.8 Application: Quad Mesh Transfer

Finally, we demonstrate a potential application of our correspondence to quad mesh transfer of artist-generated quad meshes to scanned meshes. In this experiment, we start from a set of 41 landmark points, and use weak landmark constraints with $\gamma = 5 \cdot 10^{-5}$.




Fig. 14. Quad mesh transfer using our computed correspondence. Left: input quad mesh, right: output quad mesh. Note the preservation of the prominent edge flows in the quad mesh, such as the fingernails and knuckles.

We choose to have a larger number of landmarks in this example to preserve the fine features such as fingernails and joints. The results are shown in Figure 14, with two views of the input and output quad meshes on the left and right, respectively. Note, that the edge flow of the output quad mesh closely follows the features of the hand, and the special structures in the input, such as the fingernails and knuckles, are nicely preserved in the output mesh. Such a high quality transfer can only be achieved if the computed map has a low conformal distortion, which leads to well preserved quad shapes.

6 CONCLUSION

We have suggested a novel correspondence method between non-isometric shapes, which considerably improves conformal distortion over existing techniques while preserving semantic features. Our approach is based on an alternating minimization of a quadratic energy, which is efficient, easy to implement, and flexible. In addition to demonstrating effectiveness on various benchmarks, we have shown applications to quad mesh transfer and shape interpolation.

Beyond its current utility, our formulation can serve as a framework for more involved energies on the differential of the map. For example, one can consider various metrics, such as non-isotropic, higher order, sparsity based and data-driven. We hope that our approach can be a stepping stone to generalizing the large body of literature on planar shape parameterization to the more general setting of shape correspondence.

ACKNOWLEDGEMENTS

D. Ezuz acknowledges funding from the Irwin and Joan Jacobs fellowship. J. Solomon acknowledges the generous support of Army Research Office grant W911NF-12-R-0011 (“Smooth Modeling of Flows on Graphs”), of National Science Foundation grant IIS-1838071 (“BIGDATA:F: Statistical and Computational Optimal Transport for Geometric Data Analysis”), from the MIT Research Support Committee, from an Amazon Research Award, from the MIT-IBM Watson AI Laboratory, and from the Skoltech-MIT Next Generation Program. M. Ben-Chen acknowledges funding from the European Research Council (ERC starting grant No. 714776 OPREP), and the Israel Science Foundation (grant No. 504/16).

REFERENCES

- Noam Aigerman, Shahar Z Kovalsky, and Yaron Lipman. 2017. Spherical orbifold Tutte embeddings. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 90.
- Noam Aigerman and Yaron Lipman. 2015. Orbifold Tutte Embeddings. *ACM Transactions on Graphics (TOG)* 34 (2015).
- Noam Aigerman and Yaron Lipman. 2016. Hyperbolic Orbifold Tutte Embeddings. *ACM Transactions on Graphics (TOG)* 35 (2016).
- Noam Aigerman, Roi Poranne, and Yaron Lipman. 2015. Seamless surface mappings. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 72.

- Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. 2011. The Wave Kernel Signature: A Quantum Mechanical Approach to Shape Analysis. In *International Conference on Computer Vision Workshops (ICCV Workshops)*. IEEE.
- Alexander M Bronstein, Michael M Bronstein, Benjamin Bustos, Umberto Castellani, Marco Crisani, Bianca Falcidieno, Leonidas J Guibas, Iasonas Kokkinos, Vittorio Murino, Maks Ovsjanikov, et al. 2010. SHREC 2010: robust feature detection and description benchmark. *Proc. 3DOR 2*, 5 (2010), 6.
- Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. 2006. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Sciences* 103, 5 (2006), 1168–1172.
- Oliver Burghard, Alexander Dieckmann, and Reinhard Klein. 2017. Embedding shapes with Green’s functions for global shape matching. *Computers & Graphics* 68 (2017), 1–10.
- Isaac Chao, Ulrich Pinkall, Patrick Sanan, and Peter Schröder. 2010. A simple geometric model for elastic deformations. *ACM Transactions on Graphics (TOG)* 29, 4 (2010), 38.
- Xiaobai Chen, Aleksey Golovinskiy, and Thomas Funkhouser. 2009. A Benchmark for 3D Mesh Segmentation. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 28, 3 (Aug. 2009).
- Xiaobai Chen, Abulhair Saparov, Bill Pang, and Thomas Funkhouser. 2012. Schelling points on 3D surface meshes. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 29.
- Trevor F Cox and Michael AA Cox. 2000. *Multidimensional scaling*. CRC Press.
- James Ells and Joseph H Sampson. 1964. Harmonic mappings of Riemannian manifolds. *American Journal of Mathematics* 86, 1 (1964), 109–160.
- Danielle Ezuz and Mirela Ben-Chen. 2017. Deblurring and Denoising of Maps between Shapes. In *Computer Graphics Forum*, Vol. 36. 165–174.
- Donald Geman and Chengda Yang. 1995. Nonlinear image recovery with half-quadratic regularization. *IEEE Transactions on Image Processing* 4, 7 (1995), 932–946.
- Daniela Giorgi, Silvia Biasotti, and Laura Paraboschi. 2007. SHREC: Shape Retrieval Contest: Watertight Models Track.
- Xianfeng Gu, Yalin Wang, Tony F Chan, Paul M Thompson, and Shing-Tung Yau. 2004. Genus zero surface conformal mapping and its application to brain surface mapping. *Trans. on Medical Imaging* 23, 8 (2004), 949–958.
- Behrend Heeren, Martin Rumpf, Peter Schröder, Max Wardetzky, and Benedikt Wirth. 2014. Exploring the Geometry of the Space of Shells. *Computer Graphics Forum* 33, 5 (2014), 247–256.
- Behrend Heeren, Martin Rumpf, Max Wardetzky, and Benedikt Wirth. 2012. Time-Discrete Geodesics in the Space of Shells. *Computer Graphics Forum* 31 (2012).
- Kai Hormann and Günther Greiner. 2000. *MIPS: An Efficient Global Parametrization Method*. Technical Report. DTIC Document.
- Ruqi Huang and Maks Ovsjanikov. 2017. Adjoint Map Representation for Shape Analysis and Matching. In *Computer Graphics Forum*, Vol. 36. 151–163.
- Hiroyasu Izeki and Shin Nayatani. 2005. Combinatorial harmonic maps and discrete-group actions on Hadamard spaces. *Geometriae Dedicata* 114, 1 (2005), 147–188.
- Evangelos Kalogerakis, Aaron Hertzmann, and Karan Singh. 2010. Learning 3D mesh segmentation and labeling. *ACM Transactions on Graphics* 29, 3 (2010).
- Vladimir G Kim, Yaron Lipman, and Thomas Funkhouser. 2011. Blended Intrinsic Maps. *ACM Transactions on Graphics (TOG)* 30 (2011).
- Artiom Kovnatsky, Michael M Bronstein, Xavier Bresson, and Pierre Vandergheynst. 2015. Functional Correspondence by Matrix Completion. In *Proceedings of Computer Vision and Pattern Recognition (CVPR)*.
- Artiom Kovnatsky, Michael M Bronstein, Alexander M Bronstein, Klaus Glashoff, and Ron Kimmel. 2013. Coupled Quasi-harmonic Bases. *Computer Graphics Forum* 32 (2013).
- Zorah Löhner, Matthias Vestner, Amit Boyarski, Or Litany, Ron Slossberg, Tal Remez, Emanuele Rodolà, Alexander M. Bronstein, Michael M. Bronstein, Ron Kimmel, and Daniel Cremers. 2017. Efficient Deformable Shape Correspondence via Kernel Matching. In *3D Vision (3DV)*.
- Bruno Lévy, Sylvain Petitjean, Nicolas Ray, and Jérôme Maillot. 2002. Least squares conformal maps for automatic texture atlas generation. In *ACM Transactions on Graphics (TOG)*, Vol. 21. ACM, 362–371.
- Manish Mandad, David Cohen-Steiner, Leif Kobbelt, Pierre Alliez, and Mathieu Desbrun. 2017. Variance-Minimizing Transport Plans for Inter-surface Mapping. *ACM Transactions on Graphics* 36 (2017), 14.
- Haggai Maron, Nadav Dym, Itay Kezurer, Shahar Kovalsky, and Yaron Lipman. 2016. Point Registration Via Efficient Convex Relaxation. *ACM Transactions on Graphics (TOG)* 35 (2016).
- Brent C Munsell, Pahal Dalal, and Song Wang. 2008. Evaluating shape correspondence for statistical shape analysis: A benchmark study. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 30, 11 (2008), 2023–2039.
- S. Nishikawa. 2000. *Variational Problems in Geometry*. AMS.
- Dorian Nogneng and Maks Ovsjanikov. 2017. Informative descriptor preservation via commutativity for shape matching. In *Computer Graphics Forum*, Vol. 36. 259–267.
- Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. 2012. Functional Maps: a Flexible Representation of Maps between Shapes. *ACM Transactions on Graphics (TOG)* 31 (2012).
- Daniele Panozzo, Ilya Baran, Olga Diamanti, and Olga Sorkine-Hornung. 2013. Weighted averages on surfaces. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 60.
- Ulrich Pinkall and Konrad Polthier. 1993. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics* 2, 1 (1993), 15–36.
- Emanuele Rodolà, Michael Moeller, and Daniel Cremers. 2015. Point-wise Map Recovery and Refinement from Functional Correspondence. In *Proceedings of Vision, Modeling and Visualization (VMV)*.
- Y Sahillioglu and Yücel Yemez. 2011. Coarse-to-Fine Combinatorial Matching for Dense Isometric Shape Correspondence. In *Computer Graphics Forum*, Vol. 30. 1461–1470.
- Matan Sela, Yonathan Aflalo, and Ron Kimmel. 2015. Computational caricaturization of surfaces. *Computer Vision and Image Understanding* 141 (2015), 1–17.
- Rui Shi, Wei Zeng, Zhengyu Su, Jian Jiang, Hanna Damasio, Zhonglin Lu, Yalin Wang, Shing-Tung Yau, and Xianfeng Gu. 2017. Hyperbolic harmonic mapping for surface registration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39, 5 (2017), 965–980.
- Alon Shtern and Ron Kimmel. 2014. Iterative Closest Spectral Kernel Maps. In *3D Vision (3DV)*. IEEE.
- Justin Solomon, Gabriel Peyré, Vladimir G. Kim, and Suvrit Sra. 2016. Entropic Metric Alignment for Correspondence Problems. *ACM Transactions on Graphics (TOG)* 35 (2016).
- Olga Sorkine and Marc Alexa. 2007. As-rigid-as-possible surface modeling. In *Computer Graphics Forum*, Vol. 4.
- Robert W Sumner and Jovan Popović. 2004. Deformation Transfer for Triangle Meshes. *ACM Transactions on Graphics (TOG)* 23 (2004).
- Gary KL Tam, Zhi-Quan Cheng, Yu-Kun Lai, Frank C Langbein, Yonghui Liu, David Marshall, Ralph R Martin, Xian-Fang Sun, and Paul L Rosin. 2013. Registration of 3D point clouds and meshes: a survey from rigid to nonrigid. *IEEE Transactions on Visualization and Computer Graphics* 19, 7 (2013), 1199–1217.
- Alex Tsui, Devin Fenton, Phong Vuong, Joel Hass, Patrice Koehl, Nina Amenta, David Coeurjolly, Charles DeCarli, and Owen Carmichael. 2013. Globally optimal cortical surface matching with exact landmark correspondence. In *Information Processing in Medical Imaging*, Vol. 23. NIH Public Access, 487.
- H. Urakawa. 1993. *Calculus of Variations and Harmonic Maps*. AMS.
- Oliver Van Kaick, Hao Zhang, Ghassan Hamarneh, and Daniel Cohen-Or. 2011. A survey on shape correspondence. In *Computer Graphics Forum*, Vol. 30. 1681–1707.
- M. Vestner, R. Litman, E. Rodolà, A. Bronstein, and D. Cremers. 2017. Product Manifold Filter: Non-Rigid Shape Correspondence via Kernel Density Estimation in the Product Space. In *Proceedings of Computer Vision and Pattern Recognition (CVPR)*.
- Christoph Von-Tycowicz, Christian Schulz, Hans-Peter Seidel, and Klaus Hildebrandt. 2015. Real-time Nonlinear Shape Interpolation. *ACM Transactions on Graphics (TOG)* 34 (2015).
- Yilun Wang, Junfeng Yang, Wotao Yin, and Yin Zhang. 2008. A new alternating minimization algorithm for total variation image reconstruction. *SIAM Journal on Imaging Sciences* 1, 3 (2008), 248–272.
- Kai Xu, Vladimir G Kim, Qixing Huang, Niloy Mitra, and Evangelos Kalogerakis. 2016. Data-driven shape analysis and processing. In *SIGGRAPH ASIA 2016 Courses*. ACM, 4.
- Yangyang Xu and Wotao Yin. 2013. A block coordinate descent method for regularized multiconvex optimization with applications to nonnegative tensor factorization and completion. *SIAM Journal on Imaging Sciences* 6, 3 (2013), 1758–1789.
- Xiaopeng Zheng, Chengfeng Wen, Na Lei, Ming Ma, and Xianfeng Gu. 2017. Surface Registration via Foliation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 938–947.
- Daniel Zoran and Yair Weiss. 2011. From learning models of natural image patches to whole image restoration. In *IEEE International Conference on Computer Vision (ICCV)*. IEEE, 479–486.

Appendix A. Reversibility

PROPOSITION 1. *Let $(M_1, g_1), (M_2, g_2)$ be two smooth compact Riemannian surfaces, with geodesic distance functions $d_{M_i}(\cdot, \cdot)$ given by the metrics g_i , respectively, and let $\phi_{12}: M_1 \rightarrow M_2, \phi_{21}: M_2 \rightarrow M_1$ be smooth maps. If there exists $\epsilon \geq 0$ such that:*

$$d_{M_1}(p_1, \phi_{21}(\phi_{12}(p_1))) \leq \epsilon, \quad \forall p_1 \in M_1,$$

then:

- (1) $\phi_{12}(p_1) = \phi_{12}(q_1) \Rightarrow d_{M_1}(p_1, q_1) \leq 2\epsilon, \quad \forall p_1, q_1 \in M_1.$
- (2) $\forall p_1 \in M_1 \exists p_2 \in M_2 \text{ s.t. } d_{M_1}(p_1, \phi_{21}(p_2)) \leq \epsilon.$

As a corollary of Proposition 1 we have that if the reversibility energy defined in Equation (6) is zero, then the maps ϕ_{12}, ϕ_{21} are both injective and surjective.

Proof. Let $p_1, q_1 \in M_1$, and set $p_2 = \phi_{12}(p_1)$, $q_2 = \phi_{12}(q_1)$. Further, set $\hat{p}_1 = \phi_{21}(p_2)$, and $\hat{q}_1 = \phi_{21}(q_2)$.

- (1) From the triangle inequality we have that $d_{M_1}(p_1, q_1) \leq d_{M_1}(p_1, \hat{q}_1) + d_{M_1}(\hat{q}_1, q_1)$. From the assumption (1) we have that $p_2 = q_2$ and therefore $\hat{p}_1 = \hat{q}_1$. Thus, we have $d_{M_1}(p_1, \hat{q}_1) = d_{M_1}(p_1, \hat{p}_1) \leq \epsilon$ and $d_{M_1}(\hat{q}_1, q_1) \leq \epsilon$, which gives the required result.
- (2) This follows trivially from the assumption of the Proposition if we set $p_2 = \phi_{12}(p_1)$.