

Elastic Correspondence between Triangle Meshes

D. Ezuz¹, B. Heeren², O. Azencot³, M. Rumpf², and M. Ben-Chen¹

¹Technion — Israel Institute of Technology

²University of Bonn, Institute for Numerical Simulation

³University of California, Los Angeles




Figure 1: Visualization of our results using texture transfer from the target shape (left) to the source shape with the initial map (middle) and our final map (right). Note the texture distortion in the initial map due to the incorrect matching of the airplane wing creases, which is alleviated using our crease-aware method.

Abstract

We propose a novel approach for shape matching between triangular meshes that, in contrast to existing methods, can match crease features. Our approach is based on a hybrid optimization scheme, that solves simultaneously for an elastic deformation of the source and its projection on the target. The elastic energy we minimize is invariant to rigid body motions, and its non-linear membrane energy component favors locally injective maps. Symmetrizing this model enables feature aligned correspondences even for non-isometric meshes. We demonstrate the advantage of our approach over state of the art methods on isometric and non-isometric datasets, where we improve the geodesic distance from the ground truth, the conformal and area distortions, and the mismatch of the mean curvature functions. Finally, we show that our computed maps are applicable for surface interpolation, consistent cross-field computation, and consistent quadrangular remeshing of a set of shapes.

CCS Concepts

• *Computing methodologies* → *Mesh models*;

1. Introduction

Shape correspondence is an important problem in Computer Graphics and geometry processing, with applications to deformation transfer, shape interpolation, and the analysis of shape collections, to name a few [VKZHC01]. Given source and target shapes, the goal is to find a matching point on the surface of the target for any point of the source. When the source and target shapes are triangle meshes with different connectivities, such maps take the *vertices* of the source to points on the *faces* of the target.

We consider the case of two manifold triangle meshes that have the same topology, but are not necessarily isometric. In such cases, the strong geometric prior of the preservation of geodesic distances

is no longer available, and there exists a huge set of smooth, valid maps between such surfaces. A high quality map should have a low *angle and area distortion*, facilitating downstream applications such as texture and deformation transfer. However, intrinsic geometric information alone is often not enough to yield a semantically correct map. An important *extrinsic* geometric property of semantically correct maps, is the correct alignment of prominent *curvature features*, such as the crease of an airplane wing (see Fig. 1). Yet, achieving both low metric distortion and crease alignment is difficult using existing techniques. Some mapping approaches are fully intrinsic [KLF11, AL16, MCSK*17], and therefore are not aware of extrinsic curvature dependent features. Alternatively, classic extrin-

sic approaches (e.g. [LSP08]) can match extrinsic features, but often focus on global rather than local distortion, and in addition are sensitive to the global orientation of the input shapes, since they optimize for the extrinsic deformation matrices.

We bridge this gap and achieve extrinsic feature matching, as well as low local distortion, while remaining parameterization-free and invariant to global rigid transformations of the input shapes. We accomplish this using a novel combination of a discrete *thin-shell* energy that is often used for shape deformation [HRS*14] with a recent projection-based, parameterization-free, optimization technique for *local distortion* of maps between triangle meshes [EBC17].

Starting from a set of sparse landmarks provided by the user, we initialize using an *intrinsic* map computed using existing methods [AL16]. Then we simultaneously optimize for an *elastic* deformation of the input shape while penalizing the distance from its projection on the target shape. Our energy is composed of a non-linear *membrane* energy that favors isometry, and a *bending* energy that is rotation invariant and promotes feature alignment. It is a novel modification of the classic discrete thin-shell energy, that is robust to the extreme deformations that may arise in a projection-based optimization framework, while remaining faithful to the physical behavior of the classic energy.

Our scheme yields high-quality, crease preserving maps between non-isometric shapes, that far surpass state-of-the-art methods on both the FAUST [BRLB14] and SHREC07 [GBP07] datasets. We show quantitative improvement of various error measures, specifically geodesic error from the dense ground truth on FAUST, and angular distortion, area distortion, mean curvature error and geodesic error from sparse ground truth on SHREC07. Further, our crease preserving maps are highly useful in downstream applications, as we demonstrate by applying our computed maps for generating consistent cross-fields [ACBCO17], for shape interpolation [HRWW12] and for consistent quadrangular remeshing of a set of shapes.

Our contributions. We propose a novel matching algorithm to compute a high quality correspondence map between two triangle meshes. The main characteristics are:

- Our algorithm combines a physical thin shell deformation model with a parameterization-free projection-based correspondence scheme.
- Our method is initialized with the output of an existing correspondence algorithm, which is smooth yet might have substantial local distortion and lacks feature alignment.
- Our output is a low-distortion and feature-aligned correspondence, which is highly effective for downstream applications such as texture transfer.

2. Related work

Shape correspondence. Shape correspondence and matching is a fundamental topic in geometry processing, with a considerable body of existing work. Recent surveys [VKZHC01, TCL*13] provide an excellent introduction to the state of the art, and we will therefore focus our literature review on shape correspondence methods that are closest to our approach. Many early

deformation-based approaches, that find, e.g., sparse [ZSCO*08] or dense [LSP08], correspondences, solve for an *extrinsic* deformation, namely the optimization variables depend on the local rotation that is applied to each face or vertex of the mesh to obtain the deformed mesh. These rotations affect the optimization, even though they have no impact on the final stretch and local distortion of the correspondence. Furthermore, working with extrinsic variables makes the result dependent on the global orientation of the input shapes in general.

On the other hand, intrinsic methods such as Blended Intrinsic Maps [KLF11], Seamless maps [APL15] or the Orbifold approaches [AL15, AL16] map to an intermediate domain, and therefore do not optimize for the true distortion of the correspondence, and in addition do not take into account the matching of extrinsic features such as creases. Our approach, in contrast, accounts for the actual distortion of the correspondence and for extrinsic features by using quantities such as edge lengths and dihedral angles, which are invariant to global rotations. Two more recent deformation-based correspondence approaches [AXZ*15, ZYL*17] target mostly man made shapes consisting of parts that can be represented using simple geometric primitives, and are therefore less appropriate for non-isometric manifold models.

Despite the advances in shape correspondence that were made in recent years, matching between non-isometric surfaces remains challenging. As the preservation of geodesic distances cannot serve as a prior, recent approaches use, e.g., optimal transport [MCSK*17], to regularize the correspondence. We compare our results with Mandad et al. on the SHREC07 [GBP07] dataset of non-isometric shapes, and show that our approach leads to better qualitative and quantitative results. Another common method to handle non-isometric shapes is to generate a *fuzzy correspondence*, such as a functional map [OCB*16]. However, these methods still require extracting a point-wise map from the result, for applications such as joint triangular remeshing. Finally, perhaps closest to our approach, is the method by Ezuz et al. [ESBC19]. It minimizes the harmonic energy of the forward and backward maps coupled with a reversibility constraint, which promotes bijective maps. While their reversibility approach is somewhat similar to ours, it requires a high dimensional Euclidean embedding, and is derived using different geometric principles. Furthermore, their approach uses a harmonic energy and is therefore faster, however it does not aim to preserve extrinsic features.

Elastic shape modeling. Physically-based elastic energy models have been widely used for computer graphics and geometry processing applications [RW14]. The classical model for elastically deformable surfaces is the shell model, originally introduced in a graphics context by Terzopoulos et al. [TPBF87], for thin, flexible materials. Grinspun et al. [GHDS03] introduced the discrete shell model in which a triangle mesh is a spatially-discrete representation of the mid-surface of a shell. The model was used for simulation of deformable materials under physical forces. Similar thin-shell energies have been used by Botsch et al. [BPGK06] for interactive shape deformation. Heeren and coworkers [HRWW12, HRS*14] have used the same physical model for time-discrete Riemannian analysis of shapes. In the direction of improving efficiency, the as-rigid-as-possible (ARAP) framework [SA07] is based on alternat-

ing minimisation over vertex positions and local rotations of an energy that measures deviation from rigidity. Von Radziewsky et al. [vRESH16] recently showed how model reduction can be used to efficiently evaluate elastic deformation models, including the discrete shell energy. This enables elastic models to be used in real-time applications.

Elastic shape correspondence. Already in [LDRS05] a non-linear elastic deformation energy between thin shells has been investigated for surface matching. In this approach, the matching problem is formulated on disc type parameter domains of the surface patches to be matched. This renders the surface matching problems as a classical elastic image registration problem and the membrane energy takes the role of the regularization energy, whereas the bending energy turns into a fidelity energy with respect to the matching of mean curvature functions on the parameter domains. The surface matching method in [WSSC11] picked up the matching energy from [LDRS05] and investigated matches as surfaces in the product space of the source and template geometry. Using relaxation methods from linear programming Windheuser et al. were able to robustly compute bijective triangle matching deformations with vertex to vertex correspondence. They do not allow for general deformations of source vertices onto the target surface, and the approach is quite heavy computationally due to the use of the product manifold. The elastic matching of volumetric shapes from the perspective of shape optimization has been investigated by Buhan et al. [dBDFN16]. Finally, Iglesias et al. [IRS18] have applied elastic energies for computing correspondences between level-set surfaces.

3. The elastic matching model

In a very general setup we consider two discrete surfaces, i.e. a *source* surface \mathcal{S} and a *target* surface $\tilde{\mathcal{S}}$ having $n, \tilde{n} \in \mathbb{N}$ vertices, respectively. In general, we assume $n \neq \tilde{n}$, hence \mathcal{S} and $\tilde{\mathcal{S}}$ do not share the same connectivity. The meshes are represented by the coordinate matrices with rows containing vertex positions, i.e. $X \in \mathbb{R}^{n,3}$ and $\tilde{X} \in \mathbb{R}^{\tilde{n},3}$, respectively. Since the connectivities of \mathcal{S} and $\tilde{\mathcal{S}}$ are supposed to be fixed throughout our algorithm, there are unique relationships $\mathcal{S} \leftrightarrow X$ and $\tilde{\mathcal{S}} \leftrightarrow \tilde{X}$, respectively, and we interchange notation if appropriate.

We aim at studying deformations Φ of the source surface \mathcal{S} that are constrained to the target surface $\tilde{\mathcal{S}}$ by means of a soft penalty. In detail, we consider two distinct situations. First, we only require Φ to be locally *injective*. This model is suitable for the matching of almost isometric shapes and, in particular, for partial matching. In a second step, we expand the model to be suitable for non-isometric matching problems. To this end, we additionally consider a reverse deformation $\tilde{\Phi}$ of $\tilde{\mathcal{S}}$ and study pairs of deformations $(\Phi, \tilde{\Phi})$, which are approximately inverses of each other.

3.1. Locally injective matching

In the locally injective matching case the inclusion

$$\Phi(\mathcal{S}) \subset \tilde{\mathcal{S}}, \quad (1)$$

should hold approximately and the deformation Φ should induce as little distortion as possible. We therefore define Φ such that X and




Figure 2: A sketch of the locally injective matching configuration for 1D simplicial meshes in \mathbb{R}^2 . The source mesh \mathcal{S} with coordinate matrix X (blue) is deformed via Φ to a mesh with coordinate matrix Y (dashed blue). The deformed vertices are projected based on the matrix P onto the target mesh $\tilde{\mathcal{S}}$ (red).

$\Phi(X)$ share the *same* connectivity, and use well established elastic deformation energies to control this distortion. In addition, to establish (1) on discrete surfaces, i.e. on their nodal positions X and \tilde{X} , we use a projection mapping from $\mathbb{R}^{n,3}$ to $\mathbb{R}^{\tilde{n},3}$ (see Fig. 2). These aspects are explained in detail in the following two paragraphs.

Deformation energy. We assume the deformation Φ to be defined on vertices - the deformed values of interior points are obtained by piecewise linear interpolation on the faces. In particular, $\Phi(X) \in \mathbb{R}^{n,3}$ and the resulting mesh has the same connectivity as \mathcal{S} . To simplify notation, we introduce an auxiliary variable $Y := \Phi(X)$ that contains the deformed vertex positions of \mathcal{S} and represents Φ uniquely. Now we build on the vast literature on physical deformation energies defined between two triangle meshes sharing the same connectivity, e.g. [TPBF87, GHDS03, BS08, FB11, HRWW12]. In these approaches the distortion induced by an *elastic deformation* is separated into two distinct contributions, i.e. *membrane distortion* and *bending distortion*. To this end, a generic elastic energy in the context of thin shell deformations can be written as

$$\mathcal{W}_{\text{def}}(X, Y) = \alpha \mathcal{W}_{\text{mem}}(X, Y) + \eta \mathcal{W}_{\text{bnd}}(X, Y) \quad (2)$$

with weights α and η . While $X \in \mathbb{R}^{n,3}$ is fixed throughout the algorithm, $Y \in \mathbb{R}^{n,3}$ is a primal variable that we optimize for. The definition of the two energy components will in particular ensure that (2) is invariant with respect to rigid deformations and isometric deformations minimize the membrane energy. More details also on the physical background on this energy will be given below in Section 4.

Projection mapping. To establish a suitable approximation of condition (1) on triangular surfaces we have to consider mappings between the involved discrete spaces. To this end, we consider a linear projection map $P: \mathbb{R}^{\tilde{n}} \rightarrow \mathbb{R}^n$ such that $P\tilde{X}$ is a projection of Y onto the target surface $\tilde{\mathcal{S}}$. In particular, the mapping P represents another degree of freedom. Then (1) can be achieved in a discrete setup by penalizing $\|P\tilde{X} - Y\|_X^2$, where the (squared) norm on the source surface is given as the lumped L^2 -norm

$$\|A\|_X^2 = \text{Tr}(A^\top M_X A) \quad (3)$$

with $M_X \in \mathbb{R}^{n,n}$ being the lumped mass matrix of \mathcal{S} (with the vertex areas on the diagonal).

The linear projection $P : \mathbb{R}^{\tilde{n}} \rightarrow \mathbb{R}^n$ can be represented as a (sparse) matrix $P \in \mathbb{R}^{n, \tilde{n}}$ which contains barycentric coordinates for triangles of \tilde{S} . In detail, the i -th row of the matrix P has at most three nontrivial entries $0 \leq P_{ij}, P_{ik}, P_{il} \leq 1$, such that (jkl) represents a triangle in \tilde{S} and $P_{ij} + P_{ik} + P_{il} = 1$. Let $\mathcal{P} \subset \mathbb{R}^{n, \tilde{n}}$ the set of all matrices fulfilling these properties which is denoted as the set of admissible maps. This definition implies that $(P\tilde{X})_i^T \in \mathbb{R}^3$ is a point on the discrete surface \tilde{S} for $i = 1, \dots, n$. In particular, $(P\tilde{X})_i^T$ will be a good approximation of the deformed position of the i -th vertex in S whenever $\|P\tilde{X} - Y\|_X$ is small. Note that we only project vertices, i.e. corresponding edges/faces might not be mapped onto the target surface.

Total energy. Altogether we obtain the following variational problem: For fixed $X \in \mathbb{R}^{n,3}$ and $\tilde{X} \in \mathbb{R}^{\tilde{n},3}$ we minimize the energy

$$\mathcal{E}_{\text{inj}}(Y, P) = \mathcal{W}_{\text{def}}(X, Y) + \beta \|P\tilde{X} - Y\|_X^2 \quad (4)$$

for $Y \in \mathbb{R}^{n,3}$ and $P \in \mathcal{P}$, where β is a penalty parameter and \mathcal{W}_{def} a generic elastic deformation energy as in (2) that will be further specified in Section 4. Note that for an optimal energy (4) the resulting deformation Φ defined by the optimal Y will be (close to) an isometry and, in particular, locally injective. Due to the matching term we expect that there are no local overfolds in the projection $P\tilde{X}$, which is close to Y in an L^2 -sense. However, Φ is not necessarily surjective which is obvious for instance for partial matching problems.

3.2. Bijective matching

In the case of strongly non-isometric surfaces we additionally favor surjectivity, i.e. we seek for one-to-one matching deformations. To this end, we expand our matching model in a two step procedure. First, we introduce a reverse deformation $\tilde{\Phi}$ of \tilde{S} and *symmetrize* our model (4). In particular, we consider an additional deformation energy related to $\tilde{\Phi}$ and establish a suitable approximation of the condition

$$\tilde{\Phi}(\tilde{S}) \subset S. \quad (5)$$

Then, in a second step, we ensure *reversibility* with additional energy terms to imply the opposite inclusions of (1) and (5), i.e. $\tilde{S} \subset \Phi(S)$ and $S \subset \tilde{\Phi}(\tilde{S})$, respectively.

Symmetry. We introduce an auxiliary variable $\tilde{Y} \in \mathbb{R}^{\tilde{n},3}$ which contains the deformed vertex positions of the target surface \tilde{S} and represents $\tilde{\Phi}$ uniquely, i.e. $\tilde{Y} = \tilde{\Phi}(\tilde{X})$ (see Fig. 3). Furthermore, we consider a projection map $Q : \mathbb{R}^n \rightarrow \mathbb{R}^{\tilde{n}}$ to represent projections of \tilde{Y} onto S . Analogously to the definition of $\mathcal{P} \subset \mathbb{R}^{n, \tilde{n}}$ above, we define a set $\mathcal{Q} \subset \mathbb{R}^{\tilde{n}, n}$ of admissible maps, which are sparse matrices containing rows with barycentric coordinates for triangles of S . To symmetrize (4) we finally consider the additional terms $\mathcal{W}_{\text{def}}(\tilde{X}, \tilde{Y})$ and $\tilde{\beta} \|QX - \tilde{Y}\|_{\tilde{X}}^2$ where $\tilde{\beta} \in \mathbb{R}$ is a suitable penalty parameter and $\|\cdot\|_{\tilde{X}}^2$ is defined analogously to (3). In particular, \tilde{Y} and Q are additional degrees of freedom in our optimization algorithm.

Reversibility. Solely symmetrizing our model the new variables \tilde{Y} and Q are not yet coupled with the variables Y and P of the original model (4). In particular, for a one-to-one matching we have to ensure suitable approximations of the inclusions $\tilde{S} \subset \Phi(S)$ and $S \subset \tilde{\Phi}(\tilde{S})$. So far, $P\tilde{X}$ is considered as the projection of Y onto \tilde{S}




Figure 3: To symmetrize the locally injective matching model sketched in Fig. 2 we consider a deformation Φ in the opposite direction, which maps \tilde{S} with coordinate matrix \tilde{X} to a mesh with coordinate matrix \tilde{Y} , (dashed red) and a matrix Q encoding the projection onto the source mesh S as additional degrees of freedom. To ensure reversibility we use the identical matrix P for the projection of the source X onto the deformed target $\Phi(\tilde{S})$ and for the projection of the deformed source Y onto the target \tilde{S} . Note that identical barycentric coordinates, i.e. P_{ij} and P_{ik} , are used for the projection of the undeformed source vertex x_i onto the deformed target surface and the projection of the deformed source vertex y_i onto the undeformed target.

and the difference is penalized in a lumped L^2 sense in the energy. Now, we assume that the same map P represents the projection of X onto $\Phi(\tilde{S})$, given by $P\tilde{Y} \in \mathbb{R}^{n,3}$ (see Fig. 3). The corresponding projection error $P\tilde{Y} - X$ is also measured in the lumped L^2 -norm (3) to define a corresponding penalty energy. Likewise, QX was defined to reflect the projection of \tilde{Y} onto S . Thus, we analogously consider the same map Q as the projection of \tilde{X} onto $\Phi(S)$ with the projection error $QY - \tilde{X}$, and a resulting penalty energy $\|QY - \tilde{X}\|_{\tilde{X}}^2$.

Total energy. Altogether we obtain the following variational problem for the one-to-one matching of strongly non-isometric surfaces: For fixed $X \in \mathbb{R}^{n,3}$ and $\tilde{X} \in \mathbb{R}^{\tilde{n},3}$ we minimize the energy

$$\begin{aligned} \mathcal{E}_{\text{bij}}(Y, \tilde{Y}; P, Q) = & \mathcal{W}_{\text{def}}(X, Y) + \beta \|P\tilde{X} - Y\|_X^2 \\ & + \mathcal{W}_{\text{def}}(\tilde{X}, \tilde{Y}) + \tilde{\beta} \|QX - \tilde{Y}\|_{\tilde{X}}^2 \\ & + \gamma \|P\tilde{Y} - X\|_X^2 + \tilde{\gamma} \|QY - \tilde{X}\|_{\tilde{X}}^2 \end{aligned} \quad (6)$$

for $Y \in \mathbb{R}^{n,3}$ and $\tilde{Y} \in \mathbb{R}^{\tilde{n},3}$ as well as $P \in \mathcal{P}$ and $Q \in \mathcal{Q}$. Here, $\beta, \tilde{\beta}, \gamma, \tilde{\gamma} \in \mathbb{R}$ are penalty parameters to be chosen appropriately.

The difference between \mathcal{E}_{inj} and \mathcal{E}_{bij} for non isometric shapes is demonstrated in Fig. 4. Evidently minimizing \mathcal{E}_{bij} leads to better results. Note that minimizing (6) implies $QP\tilde{X} \approx \tilde{X}$ and $PQX \approx X$. Indeed, the triangle inequality yields that $\|QP\tilde{X} - \tilde{X}\|_{\tilde{X}}$ and $\|PQX - X\|_X$ become arbitrary small for large enough penalty parameters $\beta, \tilde{\beta}, \gamma$, and $\tilde{\gamma}$.

4. Elastic deformation energy

In this section we focus on the deformation energy associated with a matching deformation Φ for two discrete surfaces sharing the *same* mesh connectivities. As above, let $X \in \mathbb{R}^{n,3}$ and $Y = \Phi(X) \in$

$\mathbb{R}^{n,3}$ be the geometries of the *undeformed* and *deformed* configuration, respectively. It will often be convenient to denote quantities living on the undeformed/reference surface with a hat – in particular, those objects are not subject to optimization since X is fixed.

As mentioned before we deploy a model of thin shell elasticity in our approach. If the triangle mesh is considered to be an approximation of a middle layer of a thin elastic material of finite thickness $\delta > 0$, the corresponding membrane energy \mathcal{W}_{mem} scales as δ , whereas isometric deformations induce a bending energy \mathcal{W}_{bnd} that scales as δ^3 . Here the non-linear membrane energy takes into account stretching and shearing and encodes a preference for isometric deformations, while the bending energy (theoretically in the limit for $\delta \rightarrow 0$ only observable for isometric deformations) compares curvature related quantities such as shape operators or mean curvatures and is in particular suitable for feature matching.

4.1. Non-linear membrane energy

Considering a simple quadratic energy functional to model membrane deformations a degenerate mesh with vanishing edge length naturally appears as a minimizer. Although a total collapse can in principle be prevented by additional forcing or boundary terms, the favoring of short edges induces a bias in the optimization. To this end, we make use of the non-linear membrane energy

$$\mathcal{W}_{\text{mem}}(X, Y) = \sum_{t \in \mathcal{T}} \hat{a}_t W(\mathcal{G}_t),$$

where \hat{a}_t denotes the area of triangle t in X and $\mathcal{G}_t = \hat{g}_t^{-1} g_t \in \mathbb{R}^{2,2}$ is the geometric distortion tensor. Here $\hat{g}_t, g_t \in \mathbb{R}^{2,2}$ are the discrete first fundamental forms of X and Y , respectively, i.e. if $e_0, e_1, e_2 \in \mathbb{R}^3$ are the edges of triangle t we have $g_t = [e_1 | -e_2]^T [e_1 | -e_2]$ which is invertible if t contains no parallel edges. The hyperelastic energy density $A \mapsto W(A)$ for $A \in \mathbb{R}^{2,2}$ is given by

$$W(A) = \frac{\mu}{2} \text{tr} A + \left[\frac{\lambda}{4} \det A - \left(\frac{\mu}{2} + \frac{\lambda}{4} \right) \log \det A - \left(\mu + \frac{\lambda}{4} \right) \right], \quad (7)$$

where $\mu, \lambda \geq 0$ are physical parameters, i.e. the Lamé coefficients of linear elasticity (see [LDRS05, HRWW12] and for mathematical details [Cia00]). Note that we choose $\lambda = \mu = 1$ in all our experiments.

Since we have $W(A) \geq 0$, $W(\mathbb{1}) = 0$ and $W_A(\mathbb{1}) = 0$, the identity map is a minimizer of the energy, and the gradient of the energy is zero at the identity. Furthermore W is invariant wrt. rigid body




Figure 4: The difference between minimizing \mathcal{E}_{inj} (left) and \mathcal{E}_{bij} (right) for non isometric shapes from SHREC07 [GBP07]. The final deformation Y is shown as a solid shape, and the target shape is rendered as wireframe. Here, \mathcal{E}_{bij} significantly improves the matching in the ear region, where \mathcal{E}_{inj} leads to unwanted artifacts.

motions, i.e., if t is deformed by means of a rigid body motion we have $\mathcal{G}_t = \mathbb{1}$ and hence $W(\mathcal{G}_t) = 0$.

The first term $\sum_{t \in \mathcal{T}} \hat{a}_t \text{tr} \mathcal{G}_t$ is sensitive to changes in edge lengths and coincides with the Dirichlet energy used in [ESBC19]. The second term penalizes variations in triangle areas. In particular, growth of area is penalized quadratically whereas shrinkage is penalized logarithmically to prevent the degeneration of faces, i.e., $W(A) \rightarrow \infty$ for $\det A \rightarrow 0$. This property is indeed crucial to prevent the collapsing of triangles by a matching deformation.

But, our algorithm relies on a good initialization given e.g. by the output of an established correspondence method (as will be discussed in Sec. 5.1). This initialization might come with degenerate faces and led to an undefined membrane energy density due to a zero argument in the log term in (7). In order to deal with these situations we slightly modify W in (7) by redefining the negative log term. In detail, we extend the function $A \mapsto -\log \det A$ via a linear profile $\log \varepsilon + \frac{x-\varepsilon}{\varepsilon}$ below a small threshold $\varepsilon > 0$ such that the compound function is still differentiable. This way $W(A)$ is well-defined even if $\det A \leq 0$.

4.2. Bending energy

As the initial choice for our bending energy one might consider the *Discrete Shells* bending energy introduced in [GHDS03], i.e.

$$\mathcal{W}_{\text{DS}}(X, Y) = \sum_{e \in \mathcal{E}} \frac{(\hat{\theta}_e - \theta_e)^2}{\hat{d}_e} \hat{l}_e^2, \quad (8)$$

where $\hat{\theta}_e, \theta_e$ denote the dihedral angle at some edge e in X and Y , respectively. If $e = t \cap t'$ we have $d_e = \frac{1}{3}(a_t + a_{t'})$, and l_e denotes the edge length of e . This energy has the smallest possible stencil (i.e. two adjacent triangles) to capture bending in a mesh and has been used extensively in the computer graphics community, e.g. [BMF03, TW06, GGWZ07, FB11, HRWW12, HSTP11]. However, when fixing t and rotating t' around the common edge e the dihedral angle θ_e jumps by 2π when the faces intersect each other (see Fig. 6). Unfortunately, this is a serious drawback when facing extreme situations, e.g. large dihedral angles and faces penetrating each other, that can appear in the initial phases of the optimization. In particular, the lack of continuity of the above energy density leads to a failure of line-search methods in the optimization. To resolve this we propose a modified version of (8) that rather penalizes deviations in the *cosine* of the dihedral angles and is continuous for all pairs of dihedral angles:

$$\mathcal{W}_{\text{bnd}}(X, Y) = \sum_{e \in \mathcal{E}} \frac{(\cos \hat{\theta}_e - \cos \theta_e)^2}{\hat{d}_e} \hat{l}_e^2. \quad (9)$$

For $e = t \cap t'$ we have $\cos \theta_e = \langle n_t, n_{t'} \rangle$ where $n_t, n_{t'} \in \mathbb{R}^3$ denote the unit triangle normals of t and t' , respectively. This energy density is periodic when rotating t' around the common edge e and has proven to be very robust even for degenerated situations (see Fig. 6). Furthermore, the evaluation of (9) is cheaper than (8), since the latter requires the costly calculation $\theta_e = \arccos \langle n_t, n_{t'} \rangle$.

This modification of the bending energy provides robustness and does not affect the physical or semantic correctness of the results




Figure 5: Comparison of the Discrete Shells (DS) bending energy (Eq. (8)) and our modified bending energy (Eq. (9)). From left to right: the target texture, the initial map (HOT), our optimized map using the DS bending energy (8), and our optimized map using the modified bending energy (9). Note that the modified energy remains faithful to the DS energy, leading to very similar results, yet with the advantage of increased robustness (see Fig. 6).

in all of the conducted experiments. This can be verified by comparing the results obtained with both energies, for a pair of meshes where the Discrete Shells bending energy does not cause numerical problems. Such an example is shown in Fig. 5.

Combining \mathcal{W}_{mem} and \mathcal{W}_{bnd} we define the total elastic energy (2) with weights $\alpha, \eta > 0$. Note that due to the scaling properties discussed above we have $\eta/\alpha \sim \delta^2$, where $\delta > 0$ is the physical thickness of the approximated thin shell. Obviously, the total elastic energy (2) is invariant with respect to rigid deformations.

5. Optimization

In this section we describe the optimization pipeline in order to solve the bijective matching problem, i.e. minimizing (6). The corresponding algorithm to optimize (4) can be considered as a special case, i.e. with certain parameters set to zero and, in particular, less degrees of freedom. The degrees of freedom (DOFs) of the energy (6) are given by $Y \in \mathbb{R}^{n,3}$ and $\tilde{Y} \in \mathbb{R}^{\tilde{n},3}$ as well as $P \in \mathcal{P}$ and $Q \in \mathcal{Q}$.

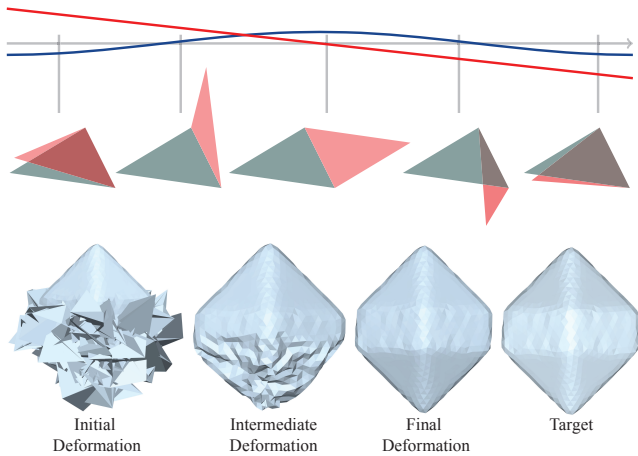


Figure 6: Robustness of bending energy. Top: different rotation angles for two adjacent triangles and corresponding graphs of discontinuous dihedral angle (red) and its periodic cosine (blue). Bottom: Reconstruction of degenerated mesh by minimizing $Y \mapsto \mathcal{W}_{\text{def}}(X, Y)$. Starting from a degenerated state (left), we first optimize with $\eta = 10^{-5}$ (middle left) and finally with $\eta = 10^{-3}$ (middle right) to restore X (right).

The numerical minimization of (6) is based on an alternating optimization strategy, i.e. we sequentially minimize one of the above matrix valued DOFs while fixing the other three. In the following, we discuss the main features of this algorithm as well as additional implementation details.

5.1. Alternating optimization

Given suitable initial mappings $P \in \mathcal{P}$ and $Q \in \mathcal{Q}$, which represent projection maps of X onto $\tilde{\mathcal{S}}$ and \tilde{X} onto \mathcal{S} , we initialize $Y := P\tilde{X}$ and $\tilde{Y} := QX$. Then we perform the alternating optimization algorithm described in Algorithm 1. Each of the K outer iterations of the algorithm consists of a sequential solution of four inner optimization problems, two optimizing for the mappings P, Q , and two optimizing for the deformations Y, \tilde{Y} .

Mapping optimization. The minimization of \mathcal{E}_{bij} with respect to P and with fixed Y, \tilde{Y}, Q is performed by solving small, constrained, linear least squares systems for each row of P , similarly to the approach suggested by Ezuz et al [EBC17]. The terms that contain P are

$$\beta \|P\tilde{X} - Y\|_X^2 + \gamma \|P\tilde{Y} - X\|_{\tilde{X}}^2,$$

thus we need to balance the quality of the projection of Y onto $\tilde{\mathcal{S}}$ with the quality of the projection of X onto $\tilde{\Phi}(\tilde{\mathcal{S}})$. As all variables except for P are fixed, this term can be written as $\|PA - B\|_X^2$, where A, B collect the fixed terms. Note that instead of considering this as the balancing of two projections in \mathbb{R}^3 , we can think of A and B as coordinates in \mathbb{R}^6 . Therefore the optimal P will give the orthogonal projection of the rows of B on a triangle mesh with the triangulation of $\tilde{\mathcal{S}}$, that is embedded in \mathbb{R}^6 and its vertex positions are given by A . This projection can be computed naively by finding for each row of B the distance to its orthogonal projection on each such 6 dimensional triangle. Let (j_0, j_1, j_2) be the triangle that realizes the smallest distance to the i -th row of B . Then, the three entries P_{ijk} for $k = 0, 1, 2$ are set to the corresponding barycentric coordinates of the closest point on the triangle. To improve efficiency, it is enough to project each point only to a relevant subset of the triangles [EBC17]. In addition, this process can be parallelized as the projection can be done in parallel for each point and triangle, and therefore it is performed on the GPU.

Deformation optimization. Minimizing \mathcal{E}_{bij} with respect to Y is done with a Quasi-Newton method while fixing all other variables.




Figure 7: Impact of bending weight. A flat, rectangular source domain $[0, 2] \times [0, 1]$ is mapped to the flat unit square (orange wireframe). We show the optimal deformation Y (blue solid, first row) as well as Y with a linear scaling in z -direction for a better visualization (blue solid, second row). The injective matching model was used with $\alpha = 10$, $\beta = 1$ (resp. $\beta = 100$ in the rightmost column) and increasing bending weights η . For the initialization of Y we added some noise in z -direction (at most 1%) to the flat source shape. The optimal deformation ranges from an isometry (left) to a uniform compression without any bending distortion (right).

Hence, we only need to evaluate the energy and the energy gradient with respect to Y . In detail, we employ a standard BFGS-method using Armijo's inexact line-search (as described in [NW99] for example) with a stopping criterion $\varepsilon = 10^{-8}$ and a maximum number of 50 BFGS iterations.

Finally, analogous strategies are used for optimizing with respect to Q and \tilde{Y} . Note that $P \in \mathcal{P}$ and $Q \in \mathcal{Q}$ by construction. For the locally injective matching problem the only degrees of freedom are P and Y , hence there are only two inner problems.

5.2. Implementation details

Global scale. Our deformation energy favors rigid body motions as solutions, where the geometric distortion tensor is the identity. We therefore scale the shapes so that the total area of each is one. This ensures that the argument of the membrane energy density (7) is not globally shifted away from its local minimum at the identity, and thus improves the robustness of our method.

Initialization. We chose to initialize our method using the Hyperbolic Orbifold Tutte Embedding (HOT) method [AL16], that pro-

duces bijective maps given a sparse set of landmarks. As HOT is intrinsic, it cannot take into consideration extrinsic features such as crease lines, which our method can align and thus significantly improve the map, as we demonstrate in the next Section.

In Fig. 8 we investigate the robustness of our method with respect to the initialization. We first modify the initial HOT map by composing it with an area-preserving map generated by the flow of a divergence free vector field. Such a transformation introduces variation in the initial map while keeping it bijective. We show the initializations and results for such a modification where the vector field is flown for shorter (VFx5) and longer (VFx10) times. The smaller distortion has no effect on our results, while the larger distortion leads to convergence to a different final map, where locally the result is smooth. Additionally, we distort the initial map by snapping each mapped point to the closest target vertex (NN). Such a distortion causes many degeneracies, therefore we use a smaller $\eta = 0.002$ to successfully handle it.




Figure 8: Effect of the initialization. We apply our method to the HOT initialization, and to three different distorted HOT initializations (see the text for the distortion details). Moderate perturbations (VFx5, NN) have no effect on our final map, whereas a strongly amplified perturbation (VFx10) leads to convergence to a different final map, that has a moderate coarse scale twist in cylindrical regions, yet is still locally smooth.

ALGORITHM 1: Alternating optimization for the bijective matching problem with $K = 300$ outer iterations and four inner variational problems to be solved sequentially via different optimization methods.

Input: Two triangle meshes $\mathcal{S}, \tilde{\mathcal{S}}$, initial $P \in \mathbb{R}^{n,\tilde{n}}$ and $Q \in \mathbb{R}^{\tilde{n},n}$
Output: Optimized P and Q

Scale $\mathcal{S}, \tilde{\mathcal{S}}$ such that each has unit total area
 Get vertex positions matrices $X \in \mathbb{R}^{n,3}$ and $\tilde{X} \in \mathbb{R}^{\tilde{n},3}$
 Initialize $Y \leftarrow P\tilde{X}$ and $\tilde{Y} \leftarrow QX$
Until (Energy increment < Threshold) or (IterationCount > K)
 $P \leftarrow \operatorname{argmin}_{P \in \mathcal{P}} \mathcal{E}_{\text{bij}}(Y, \tilde{Y}; P, Q)$ (projection)
 $Y \leftarrow \operatorname{argmin}_{Y \in \mathbb{R}^{n,3}} \mathcal{E}_{\text{bij}}(Y, \tilde{Y}; P, Q)$ (BFGS)
 $Q \leftarrow \operatorname{argmin}_{Q \in \mathcal{Q}} \mathcal{E}_{\text{bij}}(Y, \tilde{Y}; P, Q)$ (projection)
 $\tilde{Y} \leftarrow \operatorname{argmin}_{\tilde{Y} \in \mathbb{R}^{\tilde{n},3}} \mathcal{E}_{\text{bij}}(Y, \tilde{Y}; P, Q)$ (BFGS)
end

Masking edge sets. While HOT is theoretically bijective, it can generate faces which have numerically zero area due to finite numerical precision. Therefore, our optimization should be robust to such occurrences of collapsed triangles during the iterations. We have already described in Section 4 how we modify the non-linear membrane energy to be well-defined for degenerate triangles. However, such a simple extension is not appropriate for the bending energy (9), where the evaluation of the energy density at an edge requires the existence of two adjacent triangle normals. Therefore, the definition (9) is slightly modified such that we only sum over *admissible* edges, where an edge is admissible if its two adjacent triangles are not degenerate. Of course, the definition of the gradient is modified accordingly. The set of admissible edges is updated after each inner iteration step of the BFGS-method. In most experiments, we observe that almost all edges are admissible after a few outer optimization iterations, and stay admissible throughout the optimization process.

Parameters. The parameter values can be tuned according to the properties of the shapes and the desired properties of the output, e.g. whether we expect the result to be close to isometric or not.

- α The weight of the non-linear membrane energy, mainly controls the area distortion. In all our experiments we used $\alpha = 20$ along with $\lambda = \mu = 1$ for the Lamé coefficients in (7).
- η The weight of the bending energy, controls the extrinsic feature alignment. For isometric shapes, a moderate value of η is sufficient to regularize the non-linear membrane energy, and prevent undesired folding. Non-isometric shapes require a larger value of η , to favor feature matching over minimizing the amount of stretch. On the other hand, if flipped, or nearly degenerate, triangles appear in the initial map, η should have a smaller value, to allow the membrane energy to recover a smooth map. Finally, η can be used to either enforce an isometric deformation or oppress any bending distortion (see Fig. 7).
- β The weight of the mismatch energy, controls the correspondence between P and Y . It is a standard penalty coefficient, and we follow the standard guideline of penalty methods [WYYZ08] that suggests increasing β throughout the optimization process. A low value of β at the beginning of the optimization enables a significant deviation from the initial map, and β should increase to ensure the final map P corresponds to the deformation Y whose elastic energy is minimized. For all the experiments we start with $\beta = 10^3$ and linearly increase until $\beta = 2 \cdot 10^5$.
- γ The weight of the reversibility energy term, controls bijectivity. It should be high for non-isometric shapes to ensure bijectivity. We use $\gamma = 2 \cdot 10^4$ in all our experiments where reversibility was employed.

We use the same parameters for the symmetrized version of the energy, i.e. $\tilde{\beta} = \beta$ and $\tilde{\gamma} = \gamma$.

Timing. We measured example timings of our method applied on shapes with 5K vertices. The average duration of a single iteration is 1.5 seconds, and the complete optimization took 7.5 minutes. These measurements are the same order of magnitude as recent shape correspondence methods. All of our experiments were executed on a desktop machine with a TITANX GPU and an Intel Core i7 processor.

5.3. Limitations

The proposed method relies on the computation of an initial correspondence map, which is allowed to contain degenerate triangles but should not have too complicated folds and self-penetrations. To this end, our method can also be seen as a post-processing step to ensure extrinsic feature alignment on top of a fair intrinsic matching result. Unfortunately, for highly complex meshes the minimization algorithm might get stuck in a local minimum far from the global optimum. Here, a multiscale approach is required, which couples the mesh resolution on the source and the target and adapts the bending energy with an appropriate notion of scale dependent curvature information. Finally, our method is oblivious to the orientation of the resulting map, as we optimize for the first fundamental form of the deformation and not the Jacobian. On the one hand this makes our approach rotation invariant, but on the other hand we cannot incorporate into the energy a penalty for inverted triangles.

6. Results

In this section we present the results of our method for datasets that contain various types of shapes. We show quantitative and qualitative results that demonstrate improvement over the state of the art. In addition we show the applicability of our method for the applications such as consistent quadrangular remeshing and shape interpolation.

6.1. Benchmark evaluation

We measure the error with respect to the ground truth using the standard cumulative error graph [KLF11], that shows the percentage of points on the source shape (y axis) whose corresponding point on the target shape is closer than a certain geodesic distance (x axis) to the ground truth correspondence. The geodesic distances are normalized by $\sqrt{\tilde{s}}$ where \tilde{s} is the total face area of \tilde{S} . When only a sparse ground truth is given, the percentage of points is relative to the number of landmarks. Similarly, we show cumulative error graphs of the conformal and area distortion. We compute the conformal distortion per triangle as defined by Hormann et al. [HG00], and subtract 2 so that the minimal conformal distortion is zero. Area distortion per triangle is computed similarly to Nadeem et al. [NSZ*17] as $\left| \log \frac{a_t}{\hat{a}_t} \right|$ where \hat{a}_t is the area of the source triangle of S and a_t is the area of the corresponding deformed triangle.

For the qualitative evaluation we use texture transfer. We compute the texture coordinates of the target shape by projecting \tilde{X} on a plane. We then transfer them to S by applying the projection map P to the texture coordinates. As we use a texture with fine details, this visualization enables the detection of local distortion, and the assessment of semantic correctness. As the texture coordinates are computed by projection on a plane, texture discontinuities indicate regions that are parallel to the projection plane. In addition we visualize the vertex normals of the target surface by color, where the RGB values correspond to the $[x, y, z]$ coordinates of the normals. Again we transfer the normals to the source shape using P . This visualization emphasizes regions of the surface where features were not mapped correctly.

FAUST Dataset. The FAUST dataset [BRLB14] contains shapes

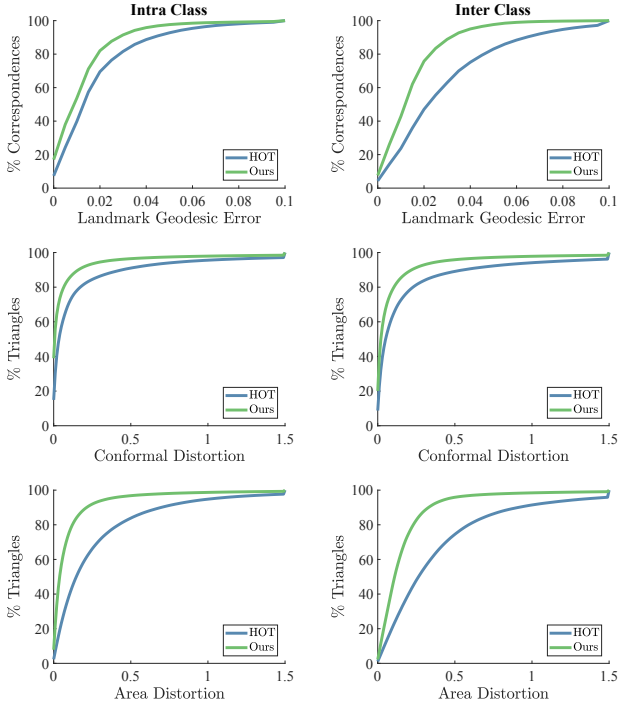


Figure 9: Quantitative results for the FAUST dataset. Our method significantly outperforms the initialization according to all distortion measures.

of 10 different humans, each of them in 10 different poses, and a ground truth correspondence is given. Since the shapes in FAUST have the same triangulation, which might bias matching algorithms, we uniformly remesh the dataset to avoid this bias. We recover the ground truth correspondence between the remeshed shapes by composing the projections of the remeshed shapes on the original corresponding shapes.

We employ the standard subsets of intra and inter classes. Note, however, that we use different parameters for the intra class and inter class, as the locally injective model (4) is sufficient for isometric shapes. For the inter class, we deploy symmetrization and reversibility, i.e. make use of the bijective model (6). We use the same values of $\alpha = 20$ and $\beta = 10^3$ initially and linearly increase until $\beta = 2 \cdot 10^5$. However, we use a larger bending weight $\eta = 0.01$ for the intra class than $\eta = 0.001$ for the inter class. The bending and reversibility terms both prevent "foldings" of the deformed surface, therefore the bending weight is larger when the reversibility term is not used. For the inter class we use $\gamma = 2 \cdot 10^4$. We use 17 landmarks as input to HOT [AL16].

The quantitative results are shown in Fig. 9. We measure the error with respect to the ground truth maps, the conformal and area distortion for each class, as specified in section 6.1. Our method significantly improves the similarity to the ground truth map, as well as the conformal and area distortion, especially for the more challenging inter class. In Fig. 10 we show the qualitative improvement due to our crease-aware method. Note especially the correct mapping of the forehead, the ears and the nose region.

SHREC'07 Dataset. We additionally test our method on 71 pairs




Figure 10: Visualization of the initial map and our results for a pair of non-isometric shapes from FAUST using texture transfer. We zoom in on the head to highlight the effect of our crease preserving map, that is semantically correct despite the distorted initialization.

of shapes from the SHREC'07 dataset [GBP07], that contains various classes that are mostly non-isometric, and the BIM benchmark [KLF11], which provides a sparse set of corresponding landmarks for each class. Since hard landmark constraints are used in HOT, we only use a subset of the landmarks given by the benchmark to evaluate the ground truth error on the remaining landmarks. For the SHREC'07 dataset we minimize the symmetric energy (6) with $\alpha = 20$, $\eta = 0.1$, $\beta = 1000$ initially and linearly increasing for the first 200 iterations and $\gamma = 2 \cdot 10^4$. These parameters are almost identical to the parameters we used for FAUST, with the exception of a larger bending weight η , which is required for the significantly non-isometric shapes of SHREC'07.

We compare our results with the initial map computed using HOT and with the shape correspondence method that was recently suggested by Manded et al. [MCSK*17] (VMTP). Additionally, we compare with two other methods that improve input correspondences, where we use the same initialization as our method (HOT): one method is based on functional maps [EBC17] (DDM) and another optimizes the reversible harmonic energy [ESBC19] (RHM). The quantitative comparison is shown in Fig. 11, and visualizations using texture mapping and normal transfer are shown in Fig. 12 and 13. Note that compared to VMTP our method yields considerably lower geodesic and conformal distortions, as well as a significantly lower area distortion compared to HOT, DDM and RHM. While RHM has the lowest conformal distortion, it does not

Class	VMTP	HOT	RHM	DDM	Ours
glasses	76.12	97.64	100.87	98.11	54.18
airplanes	263.70	175.16	183.05	188.63	105.30
ants	116.07	136.63	248.56	713.89	59.44
tables	290.57	225.59	282.18	237.05	141.57
teddies	58.76	85.26	119.68	81.44	51.90
hands	55.94	122.39	142.93	129.37	89.49
pliers	24.24	36.11	41.43	38.76	20.27
fish	59.62	87.69	84.91	89.55	48.48
birds	166.88	199.50	158.34	201.23	84.32

Table 1: Mean curvature error for each class of SHREC'07 [GBP07]. Our method produces the lowest error for all the classes except for the "hands" class.




Figure 11: Quantitative results for the SHREC07 dataset. Our results have a significantly lower geodesic and conformal distortions compared to VMTP [MCSK*17], and a significantly lower area distortion compared to HOT [AL16], DDM [EBC17] and RHM [ESBC19].




Figure 12: Visualization by texture transfer of our results compared to HOT [AL16], VMTP [MCSK*17], DDM [EBC17] and RHM [ESBC19] for a pair from each class we used from the SHREC07 dataset. See the text for details.




Figure 13: Visualization by normal transfer of our results compared to HOT [AL16], VMTP [MCSK*17], DDM [EBC17] and RHM [ESBC19] for a pair from each class we used from the SHREC07 dataset. See the text for details.

align extrinsic features, as is especially notable in the “tables” and “glasses” examples. Thus our approach yields maps which are semantic, and have both low area and conformal distortions.

Additionally, we compare the error of mean curvature alignment. Semantic maps are expected to align features, such that corresponding points should have similar mean curvature values. We compute the mean curvature per vertex of the source and target shapes using [KSNS07], and compute the error per mesh as $\|P\tilde{H} - H\|_X$ where H, \tilde{H} denote the vector with mean curvature values per vertex of the source and target shapes respectively. We sum this error across each class we used, and compare the results in Table 1. Our

mean curvature error is the lowest for almost all classes, with the exception of the “hands” class, where we believe that more landmarks are needed for a better initialization and results.

The advantage of matching mean curvature values is additionally evident in Fig. 13. The texture correspondence (left) visualizes our correct mapping of the table crease, the airplane and bird wings, and the sunglasses top crease. The normal mapping (right) serves as an additional visualization of the advantage of our approach. Note that using our result, the pushed-forward normals from the target shape match the creases of the source shape (see e.g., the handles

of the plier, the right leg of the teddy, and the bottom crease of the sunglasses).

More results are shown in Fig. 14, where we used our method (with $\alpha = 20$, $\eta = 0.002$, $\beta = 1000$ initially and linearly increasing for the first 200 iterations, $\gamma = 2 \cdot 10^4$) to compute correspondence between the left human shape and other human shapes from FAUST, and between the Giraffe and other quadrupeds from SHREC'07. The transferred texture shows that our results are highly accurate even in the case of significant non-isometric deformations.




Figure 14: Visualization of our correspondence between various non-isometric shapes from the FAUST and SHREC'07 datasets, using texture transfer.

6.2. Partial matching

We demonstrate our method for partial matching, using a simple test case, as shown in Fig. 15. We map a part of the tail of an airplane to the whole airplane shape. An initial distorted map is computed by rotation of the partial shape and projection on the target surface. While this initialization method is not robust generally, we use it as a proof of concept to show that our method can be extended to partial matching. Slight adjustments needed to be made for the partial case. First, rescaling both shapes to have the same total area is not suitable for partial matching, as it would cause the desired map to have a high area distortion which our method penalizes. In this case we do not scale the shapes. Second, since the total area of the shapes is different, the parameter values also depend on the total area of the shapes; α is divided by s the total area of \mathcal{S} , and β is divided by $s \cdot \tilde{s}$ the multiplication of the total areas of $\mathcal{S}, \tilde{\mathcal{S}}$ to




Figure 15: Demonstration of our method for partial matching of a part of the tail of an airplane to the whole airplane shape. An initial distorted map (middle) was computed by rotation of the partial shape and projection on the target surface. Our method successfully computed a map (right) that aligns the features correctly.

ensure robustness across shapes in different scales. Indeed, using this normalization, we used the same $\alpha = 20$, $\eta = 0.1$, $\beta = 1000$ for the partial matching example as we used for the full matching of shapes from SHREC07.

6.3. Applications

Consistent cross-field design. High-quality feature aligned correspondences are especially important in the context of computing consistent cross fields for a pair of shapes. Consistent cross fields are used for generating (approximately) consistent quad meshes, and are therefore required to align to the principal curvature directions of the shape. If the correspondence does not align correctly the creases of the shapes, a solution where both cross fields are aligned with the curvature directions, and also correspond under the map, will not exist. We show in Fig. 16 the quadrangulations resulting from employing the method of Azencot et al. [ACBCO17], where we either use maps computed with HOT (top row) or with our method (bottom row). Notice that HOT maps induce a significant stretch on one of the plier legs (shown both for the front and back views in the top row). In comparison, our result does not have such a shear, and, in addition, the quad edges are aligned with the curvature directions.

Consistent quadrangulation. Ultimately, shape correspondence can be used to consistently remesh a collection of shapes, for applications such as shape blending and shape analysis. To demonstrate that such an application is feasible with our mapping approach, we used our results to map a curvature-aligned quadrangulation of a source shape, shown in Fig 17 (left), to a set of other shapes, shown in Fig 17 (right). Note that the resulting quadrangulations are similarly curvature aligned, although we have not explicitly computed them as such.

Shape interpolation. We demonstrate the usage of our method for shape interpolation. As shape interpolation methods mostly require consistent triangulation of the source and target shapes, a correspondence between them is required. We remesh the target shape according to a map P by setting the new vertex positions to $P\tilde{X}$ and using the triangulation of the source shape. Then we use [HRS*14] to interpolate between the shapes. The results are shown in Fig. 18, and the supplementary video. Note the smooth interpolation of the shapes, which indicates the semantically correct mapping of features.

7. Conclusions and future work

We presented a novel elastic matching model for computing a correspondence between triangular meshes, that minimizes the metric distortion, puts crease features in correspondence and is globally rotation invariant. We optimized simultaneously for an elastic deformation of the source shape, using a modified robust elastic energy, and the distance of its projection from the target shape with an alternating minimization approach. Our method is applicable to both isometric and non-isometric shapes, as we demonstrated by achieving results which exceed the state of the art on two shape datasets as measured by multiple metrics. Finally, we showed that our approach yields high quality maps which are applicable to challenging use cases such as shape interpolation, joint cross field design and joint quadrangulation.

Our projection matrix can be seen as a *functional map*, defined in a piecewise linear basis, and it is interesting to explore the applicability of spectral bases in the context of elastic matching. Furthermore, as crease features are scale dependent, a promising avenue of future work is to generalize the method to use a multiscale hierarchical approach.

Acknowledgements

D. Ezuz acknowledges funding from the Irwin and Joan Jacobs fellowship. B. Heeren and M. Rumpf acknowledge support of the Hausdorff Center for Mathematics. B. Heeren M. Rumpf and M. Ben Chen acknowledge the support by a grant from the GIF, the German-Israeli Foundation for Scientific Research and Development (grant number I-1339-407.6/2016). O. Azencot acknowledges the supported by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 793800, a Zuckerman STEM Leadership Postdoctoral Fellowship. M. Ben Chen acknowledges the support of the Israel Science Foundation (grant No. 504/16).

References

[ACBCO17] AZENCOT O., CORMAN E., BEN-CHEN M., OVSJANIKOV M.: Consistent functional cross field design for mesh quadrangulation.




Figure 16: We compute consistent cross fields to quadrangulate a pair of shapes and we provide views from the front and the back. Using maps computed with HOT (top row) yields a significant shear on the plier leg. This shear is not observed when our correspondences are used (bottom row).




Figure 17: Consistent quadrangulation of a set of shapes by computing a curvature aligned quadrangulation of one shape (left) and mapping it to the other shapes (right) using our computed map. Note that the mapped quadrangulations are curvature aligned, although they were not computed as such explicitly.

ACM Trans. Graph. 36, 4 (2017), 92:1–92:13. 2, 12

[AL15] AIGERMAN N., LIPMAN Y.: Orbifold Tutte embeddings. *ACM Trans. Graph.* 34, 6 (2015), 190:1–190:12. 2

[AL16] AIGERMAN N., LIPMAN Y.: Hyperbolic orbifold Tutte embeddings. *ACM Trans. Graph.* 35, 6 (2016), 217:1–217:14. 1, 2, 7, 9, 10, 11

[APL15] AIGERMAN N., PORANNE R., LIPMAN Y.: Seamless surface mappings. *ACM Trans. Graph.* 34, 4 (2015), 34:1–34:10. 2

[AXZ*15] ALHASHIM I., XU K., ZHUANG Y., CAO J., SIMARI P., ZHANG H.: Deformation-driven topology-varying 3D shape correspondence. *ACM Trans. Graph.* 34, 6 (2015), 236:1–236:13. 2

[BMF03] BRIDSON R., MARINO S., FEDKIW R.: Simulation of clothing with folds and wrinkles. In *Proc. ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2003), pp. 28–36. 5

[BPGK06] BOTSCH M., PAULY M., GROSS M. H., KOBELT L.: PriMo: coupled prisms for intuitive surface modeling. In *Proc. Eurographics Symposium on Geometry Processing* (2006), pp. 11–20. 2

[BRLB14] BOGO F., ROMERO J., LOPER M., BLACK M. J.: FAUST: Dataset and evaluation for 3D mesh registration. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition* (2014), pp. 3794–3801. 2, 8

[BS08] BOTSCH M., SORKINE O.: On linear variational surface deformation methods. *IEEE Trans. Vis. Comput. Graph.* 14, 1 (2008), 213–230. 3

[Cia00] CIARLET P. G.: *Mathematical elasticity. Vol. III*, vol. 29 of *Studies in Mathematics and its Applications*. North-Holland Publishing Co., Amsterdam, 2000. Theory of shells. 5

[dBDFN16] DE BUHAN M., DAPOGNY C., FREY P., NARDONI C.: An optimization method for elastic shape matching. *C. R. Math. Acad. Sci. Paris* 354, 8 (2016), 783–787. 3

[EBC17] EZUZ D., BEN-CHEN M.: Deblurring and denoising of maps between shapes. *Comput. Graph. Forum* 36, 5 (2017), 165–174. 2, 6, 9, 10, 11

[ESBC19] EZUZ D., SOLOMON J., BEN-CHEN M.: Reversible harmonic maps between discrete surfaces. *ACM Trans. Graph.*, to appear (2019). 2, 5, 9, 10, 11

[FB11] FRÖHLICH S., BOTSCH M.: Example-driven deformations based on discrete shells. *Comput. Graph. Forum* 30, 8 (2011), 2246–2257. 3, 5

[GBP07] GIORGI D., BIASOTTI S., PARABOSCHI L.: SHREC: Shape retrieval contest: Watertight models track, 2007. 2, 5, 9

[GGWZ07] GARG A., GRINSUN E., WARDETSKY M., ZORIN D.: Cubic shells. In *Proc. ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2007), pp. 91–98. 5

[GHDS03] GRINSUN E., HIRANI A. N., DESBRUN M., SCHRÖDER P.: Discrete shells. In *Proc. ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2003), pp. 62–67. 2, 3, 5




Figure 18: Shape interpolation using [HRS*14]. The target shape is remeshed using our result (right) so that the triangulation of the source (left) and target shapes correspond. The resulting interpolation (three middle shapes) is smooth, which indicates a correct semantic mapping of the airplane's features.

- [HG00] HORMANN K., GREINER G.: MIPS: An efficient global parametrization method, 2000. 8
- [HRS*14] HEEREN B., RUMPF M., SCHRÖDER P., WARDETZKY M., WIRTH B.: Exploring the geometry of the space of shells. *Comput. Graph. Forum* 33, 5 (2014), 247–256. 2, 12, 14
- [HRWW12] HEEREN B., RUMPF M., WARDETZKY M., WIRTH B.: Time-discrete geodesics in the space of shells. *Comput. Graph. Forum* 31, 5 (2012), 1755–1764. 2, 3, 5
- [HSTP11] HILDEBRANDT K., SCHULZ C., TYCOWICZ C. V., POLTHIER K.: Interactive surface modeling using modal analysis. *ACM Trans. Graph.* 30, 5 (2011), 119:1–11. 5
- [IRS18] IGLESIAS J. A., RUMPF M., SCHERZER O.: Shape-aware matching of implicit surfaces based on thin shell energies. *Found. Comput. Math.* 18, 4 (2018), 891–927. 3
- [KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T.: Blended intrinsic maps. *ACM Trans. Graph.* 30, 4 (2011), 79:1–79:12. 1, 2, 8, 9
- [KSNS07] KALOGERAKIS E., SIMARI P., NOWROUZEZAHRAI D., SINGH K.: Robust statistical estimation of curvature on discretized surfaces. In *Proc. Eurographics Symposium on Geometry Processing* (2007). 11
- [LDRS05] LITKE N., DROSKE M., RUMPF M., SCHRÖDER P.: An image processing approach to surface matching. In *Proc. Eurographics Symposium on Geometry Processing* (2005), pp. 207–216. 3, 5
- [LSP08] LI H., SUMNER R. W., PAULY M.: Global correspondence optimization for non-rigid registration of depth scans. *Comput. Graph. Forum* 27, 5 (2008), 1421–1430. 2
- [MCSK*17] MANDAD M., COHEN-STEINER D., KOBELT L., ALLIEZ P., DESBRUN M.: Variance-minimizing transport plans for inter-surface mapping. *ACM Trans. Graph.* 36, 4 (2017), 39:1–39:14. 1, 2, 9, 10, 11
- [NSZ*17] NADEEM S., SU Z., ZENG W., KAUFMAN A., GU X.: Spherical parameterization balancing angle and area distortions. *IEEE Trans. Vis. Comput. Graph.* 23, 6 (2017), 1663–1676. 8
- [NW99] NOCEDAL J., WRIGHT S. J.: *Numerical Optimization*. Springer Series in Operations Research. Springer-Verlag, New York, 1999. 7
- [OCB*16] OVSJANIKOV M., CORMAN E., BRONSTEIN M., RODOLÀ E., BEN-CHEN M., GUIBAS L., CHAZAL F., BRONSTEIN A.: Computing and processing correspondences with functional maps. In *SIGGRAPH ASIA 2016 Courses* (2016), ACM, p. 9. 2
- [RW14] RUMPF M., WARDETZKY M.: Geometry processing from an elastic perspective. *GAMM-Mitteilungen* 37, 2 (2014), 184–216. 2
- [SA07] SORKINE O., ALEXA M.: As-rigid-as-possible surface modeling. In *Proc. Eurographics Symposium on Geometry Processing* (2007), pp. 109–116. 2
- [TCL*13] TAM G. K., CHENG Z.-Q., LAI Y.-K., LANGBEIN F. C., LIU Y., MARSHALL D., MARTIN R. R., SUN X.-F., ROSIN P. L.: Registration of 3d point clouds and meshes: a survey from rigid to nonrigid. *IEEE Trans. Vis. Comput. Graph.* 19, 7 (2013), 1199–1217. 2
- [TPBF87] TERZOPOULOS D., PLATT J., BARR A., FLEISCHER K.: Elastically deformable models. In *Proc. SIGGRAPH* (1987), vol. 21, pp. 205–214. 2, 3
- [TW06] THOMASZEWSK B., WACKER M.: Bending models for thin flexible objects. In *WSCG Short Comm.* (2006). 5
- [VKZHC011] VAN KAICK O., ZHANG H., HAMARNEH G., COHEN-OR D.: A survey on shape correspondence. *Comput. Graph. Forum* 30, 6 (2011), 1681–1707. 1, 2
- [vRESH16] VON RADZIEWSKY P., EISEMANN E., SEIDEL H.-P., HILDEBRANDT K.: Optimized subspaces for deformation-based modeling and shape interpolation. *Computers & Graphics* 58 (2016), 128–138. 3
- [WSSC11] WINDHEUSER T., SCHLICKWEI U., SCHMIDT F. R., CREMERS D.: Geometrically consistent elastic matching of 3d shapes: A linear programming solution. In *Proc. IEEE International Conference on Computer Vision* (2011), pp. 2134–2141. 3
- [WYYZ08] WANG Y., YANG J., YIN W., ZHANG Y.: A new alternating minimization algorithm for total variation image reconstruction. *SIAM J. Imaging Sci.* 1, 3 (2008), 248–272. 8
- [ZSCO*08] ZHANG H., SHEFFER A., COHEN-OR D., ZHOU Q., VAN KAICK O., TAGLIASACCHI A.: Deformation-driven shape correspondence. *Comput. Graph. Forum* 27, 5 (2008), 1431–1439. 2
- [ZYL*17] ZHU C., YI R., LIRA W., ALHASHIM I., XU K., ZHANG H.: Deformation-driven shape correspondence via shape recognition. *ACM Trans. Graph.* 36, 4 (2017), 51:1–51:12. 2