

Geometrically Consistent Elastic Matching of 3D Shapes: A Linear Programming Solution

Thomas Windheuser¹, Ulrich Schlickewei¹, Frank R. Schmidt² and Daniel Cremers¹

¹Technische Universität München

²University of Western Ontario

Abstract

We propose a novel method for computing a geometrically consistent and spatially dense matching between two 3D shapes. Rather than mapping points to points we match infinitesimal surface patches while preserving the geometric structures. In this spirit we consider matchings as diffeomorphisms between the objects' surfaces which are by definition geometrically consistent. Based on the observation that such diffeomorphisms can be represented as closed and continuous surfaces in the product space of the two shapes we are led to a minimal surface problem in this product space. The proposed discrete formulation describes the search space with linear constraints. Computationally, our approach leads to a binary linear program whose relaxed version can be solved efficiently in a globally optimal manner. As cost function for matching, we consider a thin shell energy, measuring the physical energy necessary to deform one shape into the other. Experimental results demonstrate that the proposed LP relaxation allows to compute high-quality matchings which reliably put into correspondence articulated 3D shapes. Moreover a quantitative evaluation shows improvements over existing works.

1. Introduction

An increasing number of digitized three-dimensional objects has become available over the last years due to the technical progress in acquisition hardware like laser scanners or medical imaging devices. Such objects originate from a variety of different domains including biology, medicine, industrial design or computer animation. This rapid growth in stored data brings about the need for reliable algorithms to organize this data. One of the cornerstone problems in this context is the matching problem: Given two three-dimensional objects, find a meaningful correspondence between the object's surfaces. To date there is no efficient and optimal algorithm for this problem.

In this work, we propose a novel framework for finding




Figure 1. We propose to cast the dense elastic matching of surfaces in 3D as a codimension-two minimal surface problem which aims at minimizing the distortion when transforming one shape into the other. We show that a consistent discretization of this minimal surface problem gives rise to an integer linear program. By means of LP relaxation we can compute near-optimal matchings such as the one shown above. These matchings are dense triangle-wise matchings. (For visualization we combined triangles to patches and colored them consistently with their corresponding patch.)

an optimal geometrically consistent matching between two surfaces. We formulate shape matching as a minimal surface problem which allows for a linear programming discretization. This model comes with a sound physical interpretation and allows to compute high-quality matching without need for initialization.

1.1. Related Work

Interestingly, in one less dimension the matching of *planar* shapes can be solved by means of dynamic programming in runtimes which are subcubic in the number of points on each shape [18]. Unfortunately the concepts of dynamic programming and Dijkstra's algorithm do not extend to the third dimension where the solution is no longer

a shortest path but a minimal closed surface in a higher-dimensional space. Therefore existing approaches for three-dimensional shape matching typically rely on *local* optimization techniques. Our approach tries to overcome this limitation by describing these closed surfaces with linear constraints involving the boundary operator. Inspired by Sullivan [20], the boundary operator was previously introduced in the context of image segmentation by Grady [9] and Schoenemann et al. [19].

The paradigm of the Gromov–Hausdorff framework, proposed by Mémoli and Sapiro in [16], is to find the correspondence which minimizes the geodesic distortion. Bronstein et al. [2] proposed an efficient method for computing such correspondences in a coarse-to-fine strategy much akin to optical flow algorithms.

Other approaches to shape matching employ techniques from conformal geometry [24, 15] or Riemannian geometry [12]. The physically motivated energy model we use in this work is related to the works of Litke et al. [14] and of Rumpf and Wirth [25].

All the above-mentioned methods have in common that they use a local optimization technique to minimize a non-convex energy. As a consequence, the quality of solutions depends heavily on a good initialization and an appropriately designed coarse-to-fine strategy. In addition, solutions do not come with any optimality guarantees, which implies that in principle they can be arbitrarily bad.

Recently, two methods with a more *global* flavor have been proposed. On the one hand, Zeng and coworkers [27] formulate shape matching as a graph matching problem of third order and apply the QPBO algorithm [22]. Although the overall approach does not guarantee globally optimal solutions, it is able to detect when a proposed matching pair is globally optimal. Two major drawbacks of this approach are that firstly it suffers from a very high computational complexity, considering all triples of possible matchings. In practice it allows only the matching of a few feature points which is then postprocessed with a local method. Secondly, this approach lacks a continuous counterpart, as it merely matches discrete points rather than surface elements.

On the other hand, Lipman and Daubechies [13] recently proposed to compare surfaces of genus zero and open surfaces using optimal mass transport and conformal geometry. Computationally, this amounts to solving a linear program in n^2 variables where n is the number of vertices used in the discretization of the surfaces. The problem with this approach is that no spatial regularity is imposed on the matchings.

1.2. Contribution

We propose a novel formulation for the shape matching problem based on finding an optimal surface of codimension 2 in the product of the two shape surfaces. This sur-

face minimizes the physical deformation energy needed for deforming one shape into the other. We derive a consistent discretization of the continuous framework and show that the discrete minimal surface problem amounts to a linear program. Compared to existing approaches the proposed framework has the following advantages:

- The LP formulation is a global approach allowing to compute matchings which are independent of initialization with no postprocessing.
- The proposed method guarantees a geometrically consistent matching in the sense that the surfaces are mapped into one another in a continuous and orientation preserving manner.
- We provide a discretization of the set of surface diffeomorphisms by means of linear constraints. This is quite remarkable because in previous formulations the diffeomorphism constraint is highly non-linear and computationally very difficult [26].
- The algorithmic formulation is independent of the particular choice of deformation energy and can be applied universally. As an example, we show that one can also incorporate local feature similarity in order to improve performance.
- Experiments demonstrate that reliable and dense matchings are obtained even for larger problem instances with no need for postprocessing.

2. From Shape Matching to Minimal Surfaces: The Continuous Setting

In this section we outline a shape matching model which is based on minimizing physically motivated energies. We then show how this problem can be translated into an equivalent problem of finding a minimal codimension-two surface in a four-dimensional space.

2.1. Shape Matching based on Minimizing Deformation Energies

In the following, we assume that the two shapes $X, Y \subset \mathbb{R}^3$ are differentiable, oriented, closed surfaces. While most 3D shape matching approaches like to interpret a matching just as a bijective mapping between the surface points of these shapes, we pursue a fundamentally different approach. The main reason is that in general, bijections do not respect the underlying two-dimensional structure of surfaces. In fact, there are even continuous bijections between a line and a two-dimensional patch like the continuous *Hilbert space-filling curve* [10]. Therefore, we propose to search for diffeomorphisms instead of bijections.

Diffeomorphisms $f : X \rightarrow Y$ are bijections for which both, f and f^{-1} are differentiable. This does not only cope

with the dimensionality problem presented above, but it also helps us to propose an energy function that is symmetric in f and its inverse. As a result, the optimal matching f between X and Y , also gives rise to the optimal matching between Y and X , namely $f^{-1} : Y \rightarrow X$.

In the following, we formulate the shape matching problem as an optimization problem over the set of orientation preserving diffeomorphisms between X and Y

$$\inf_{f \in \text{Diff}^+(X, Y)} E(f) + E(f^{-1}) \quad (1)$$

where E is a suitable energy on the class of all diffeomorphisms between surfaces and $\text{Diff}^+(X, Y)$ is the set of orientation preserving diffeomorphisms between X and Y . Note that we choose a symmetric problem formulation, penalizing at the same time deformation energy of X into Y and of Y into X . This is necessary because usually E takes different values on f and on f^{-1} .

The energy functional we use is borrowed from elasticity theory in physics [3]. Interpret the surfaces X and Y as ‘‘thin shells’’. Now we try to find the deformation of X into Y which requires the least stretching and bending energy. Such models usually consist of a membrane energy E_{mem} and a bending energy E_{bend} penalizing deformations in the first and in the second fundamental forms of the surfaces. In this work we use the following formulation:

$$E(f) = \underbrace{\int_X (\text{tr}_{g_X} \mathbf{E}) + \mu \text{tr}_{g_X} (\mathbf{E}^2)}_{E_{\text{mem}}} + \lambda \underbrace{\int_X (H_X(x) - H_Y(f(x)))^2}_{E_{\text{bend}}} \quad (2)$$

where $\mathbf{E} = f^*g_Y - g_X$ is the difference between the metric tensors of X and Y , typically called the *Lagrange strain tensor*, $\text{tr}_{g_X}(\mathbf{E})$ is the norm of this tensor (see [6]), H_X and H_Y denote the mean curvatures and μ and λ are parameters which determine the elasticity and the bending property of the material. This energy is a slightly simplified version of Koiter’s thin shell energy [11].

2.2. Diffeomorphisms and their Graph Surfaces

Given an orientation preserving diffeomorphism $f : X \rightarrow Y$ we obtain a surface $\Gamma \subset X \times Y$ in the Euclidean product of X and Y by passing to the graph

$$\Gamma = \{(x, f(x)) \mid x \in X\} \subset X \times Y. \quad (3)$$

The surface Γ comes with two natural projections $\pi_X : \Gamma \rightarrow X$, $(x, f(x)) \mapsto x$ and $\pi_Y : \Gamma \rightarrow Y$, $(x, f(x)) \mapsto f(x)$. A diffeomorphism is completely characterized by its graph:

Proposition 1 (graph surfaces). *Let Γ be the graph of a diffeomorphism $f : X \rightarrow Y$. Then*

- (i) Γ is a differentiable, connected, closed surface in the product space $X \times Y$.
- (ii) The projections π_X and π_Y are both diffeomorphisms.
- (iii) The two orientations which Γ naturally inherits from X and Y coincide.

*Vice versa, any surface $\Gamma \subset X \times Y$ which satisfies (i), (ii) and (iii) is the graph of an orientation-preserving diffeomorphism between X and Y . We call such surfaces **graph surfaces**.*

The energy $E(f)$ can be expressed as

$$E(f) = \tilde{E}(\Gamma) \quad (4)$$

where $\tilde{E}(\Gamma) = E(\pi_Y \circ (\pi_X)^{-1}) + E(\pi_X \circ (\pi_Y)^{-1})$.

The outcome of the above discussion is that the optimization problem (1) can be phrased as an optimization problem over the set of surfaces in $X \times Y$ which then reads

$$\begin{aligned} & \inf \tilde{E}(\Gamma) \\ & \text{subject to } \Gamma \subset X \times Y \text{ is a graph surface} \end{aligned} \quad (5)$$

We remark that the idea of casting optimal diffeomorphism problems as minimal surface problems has been applied previously in the theory of nonlinear elasticity [8]. In the setup of shape matching, it is related to the approach that Tagare [21] proposed for the matching of 2D shapes. It was reformulated as an orientation preserving diffeomorphism approach in [18].

3. The Discrete Setting

In this section we develop a discrete counterpart of the notion of graph surfaces in $X \times Y$ introduced in Section 2.2. We start in 3.1 with the definition of discrete surface patches in $X \times Y$. These patches will be the building blocks for discrete graph surfaces introduced in 3.2. Finally in 3.3 we give a discrete version of the energy.

3.1. Discrete Surface Patches

Let $X = (V_X, E_X, F_X)$ be a triangulated oriented surface mesh, consisting of a set of vertices V_X , of directed edges E_X and of oriented triangles F_X .

A priori, edges on X do not have a preferable orientation. Therefore, we fix an orientation for each edge on X . Thus, whenever two vertices a_1 and a_2 of X are connected by an edge, either $\binom{a_1}{a_2} \in E_X$ or $\binom{a_2}{a_1} = -\binom{a_1}{a_2} \in E_X$. We extend the set of edges by *degenerate edges*

$$\overline{E}_X = E_X \cup \left\{ \binom{a}{a} \mid a \in V_X \right\}. \quad (6)$$

By assumption, the triangular faces of X are oriented. If the vertices a_1, a_2, a_3 build an oriented triangle on X , then

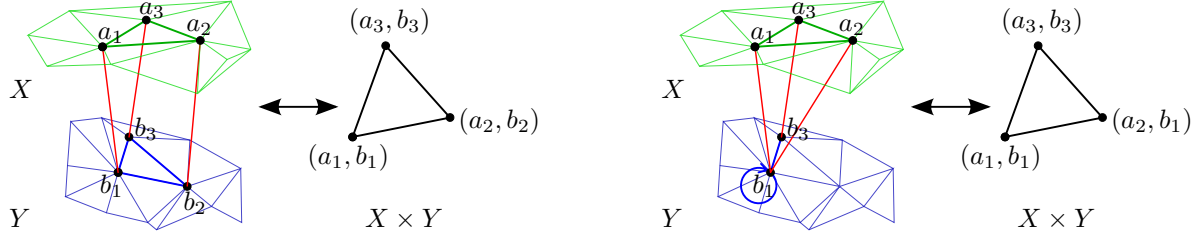


Figure 2. To assure a geometrically consistent, elastic matching from mesh X to mesh Y , we define a space of feasible solutions which is spanned by a set of 45 basic matchings among triangles, edges and vertices on either mesh. Two representative matchings and their corresponding representation in the product space $X \times Y$ are shown. **Left image:** The triangle $(a_1, a_2, a_3)^T$ on surface X is matched to triangle $(b_1, b_2, b_3)^T$ on Y by assigning vertex a_i to vertex b_i . This directly corresponds to the triangle with vertices (a_i, b_i) in the product graph. **Right image:** The triangle $(a_1, a_2, a_3)^T$ is matched to the edge $(b_1, b_3)^T$, represented here as degenerate triangle $(b_1, b_1, b_3)^T$.

$\begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix} = \begin{pmatrix} a_2 \\ a_3 \\ a_1 \end{pmatrix} = \begin{pmatrix} a_3 \\ a_1 \\ a_2 \end{pmatrix} \in F_X$. As for the edges, we extend the set of triangles by *degenerate triangles*

$$\bar{F}_X = F_X \cup \left\{ \begin{pmatrix} a_1 \\ a_2 \\ a_2 \end{pmatrix} \mid a_1, a_2 \in V_X, \pm \begin{pmatrix} a_1 \\ a_2 \end{pmatrix} \in \bar{E}_X \right\}. \quad (7)$$

Notice that degenerate triangles can consist of only one or of two vertices. The existence of these degenerate triangles will allow stretching or compression of parts of the surface as becomes apparent below (see also Figure 7).

Next, we introduce product surface triangles for two triangular meshes X and Y . Define the product of X and Y by the set of vertices $V = V_X \times V_Y$, the set of edges $E = \bar{E}_X \times \bar{E}_Y$ and the set of product triangles

$$F := \left\{ \begin{pmatrix} a_1, b_1 \\ a_2, b_2 \\ a_3, b_3 \end{pmatrix} \mid \begin{array}{l} f_1 = \begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix} \in \bar{F}_X, \\ f_2 = \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix} \in \bar{F}_Y, \\ f_1 \text{ or } f_2 \text{ non-degenerate} \end{array} \right\} \quad (8)$$

The product triangles in F are the basic pieces which are later glued to discrete graph surfaces. For shape matching, a product triangle

$$\begin{pmatrix} a_1, b_1 \\ a_2, b_2 \\ a_3, b_3 \end{pmatrix} \in F \quad (9)$$

is interpreted as setting vertex $a_i \in V_X$ in correspondence with vertex $b_i \in V_Y$.

Given two non-degenerate triangles $a = \begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix} \in F_X$ and $b = \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix} \in F_Y$ we allow for 45 different matchings between them: 3 orientation-preserving bijective matchings, 36 triangle-to-edge matchings and 6 triangle-to-vertex matchings. These degenerate matchings allow us to handle infinitesimal stretching and compression in the proposed framework. A visualization of two of the 45 possibilities is given in Figure 2.

3.2. Discrete Surfaces

Following Section 2 a diffeomorphism can be represented as a surface $\Gamma \subset X \times Y$ satisfying conditions (i), (ii) and (iii). In this section we derive discrete versions of these properties.

Definition 2. A discrete surface in $X \times Y$ is a subset $\Gamma \subset F$. The set of all discrete surfaces is denoted by $\text{surf}(X \times Y)$.

As we have seen above, a product triangle in F can be interpreted as matching a triangle on X to a triangle on Y . Thus, the intuitive meaning of a discrete surface $\Gamma \subset F$ is a set of correspondences between triangles on X and Y . Imposing the discrete counterparts of (i), (ii) and (iii) on such a discrete surface will result in a geometrically consistent, bijective matching.

Discrete version of (i): In the following we will find a condition which guarantees the continuity of our matching.

Recall that the boundary operator for triangle meshes [5] maps triangles to their oriented boundary. We extend this definition to the product graph G .

As for the sets E_X and E_Y we choose arbitrary orientations for each product edge $e \in E$. By means of these orientations we define for any edge $\begin{pmatrix} v_1 \\ v_2 \end{pmatrix}$ connecting two vertices $v_1, v_2 \in V$ a vector in $O\left(\begin{pmatrix} v_1 \\ v_2 \end{pmatrix}\right) \in \mathbb{Z}^{|E|}$ whose e -th entry is given by

$$O\left(\begin{pmatrix} v_1 \\ v_2 \end{pmatrix}\right)_e = \begin{cases} 1 & \text{if } e = \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} \\ -1 & \text{if } e = \begin{pmatrix} v_2 \\ v_1 \end{pmatrix} \\ 0 & \text{else.} \end{cases} \quad (10)$$

The triangles in F naturally inherit orientations from the triangles in F_X and F_Y . This allows us to define the boundary operator as follows.

Definition 3. The boundary operator $\partial : F \rightarrow \mathbb{Z}^{|E|}$ is de-

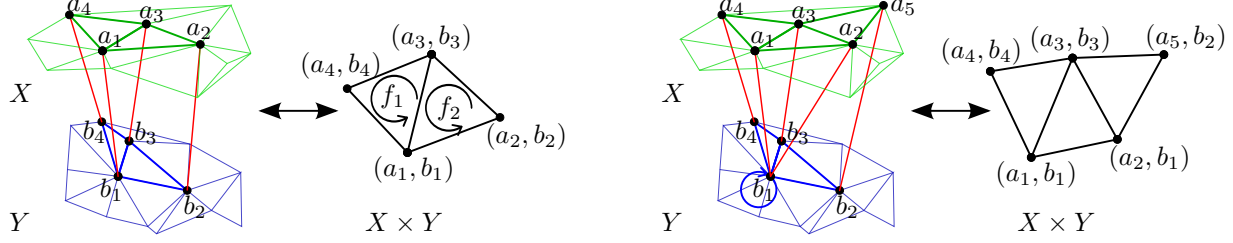


Figure 3. To ensure that neighboring triangles on X are matched with neighboring triangles on Y , we impose the closeness condition. **Left image (general case):** The triangles $(a_1, a_2, a_3)^T$ and $(b_1, b_2, b_3)^T$ are matched resulting in activating f_2 . The boundary condition $\partial\Gamma = 0$ ensures that the matching continues with a correspondence whose triangles in X and Y are positively incident to $(a_1, a_3)^T$ and $(b_1, b_3)^T$ respectively. This constraint is satisfied for example by triangle f_1 which is visualized here. **Right image (stretching):** The stretching is achieved by matching triangle $(a_1, a_2, a_3)^T$ to edge $(b_3, b_1)^T$. Again, the geometric consistency is granted by the boundary condition evaluated on the product edges $((a_2, b_1), (a_3, b_3))^T$ and $((a_3, b_3), (a_1, b_1))^T$.

fined by

$$\partial \begin{pmatrix} a_1, b_1 \\ a_2, b_2 \\ a_3, b_3 \end{pmatrix} := O \begin{pmatrix} a_1, b_1 \\ a_2, b_2 \end{pmatrix} + O \begin{pmatrix} a_2, b_2 \\ a_3, b_3 \end{pmatrix} + O \begin{pmatrix} a_3, b_3 \\ a_1, b_1 \end{pmatrix}, \quad (11)$$

where the $a_i \in V_X$ and $b_i \in V_Y$ form triangles on X resp. on Y and $\begin{pmatrix} a_i, b_i \\ a_j, b_j \end{pmatrix}$ is the product edge connecting the product vertices (a_i, b_i) and (a_j, b_j) .

The boundary operator is linearly extended to a map

$$\partial : \text{surf}(X \times Y) \rightarrow \mathbb{Z}^{|E|}. \quad (12)$$

A discrete surface Γ in $X \times Y$ is **closed** if $\partial\Gamma = 0$.

The closeness condition ensures that adjacent triangles on X are in correspondence with adjacent triangles on Y and therefore guarantees the geometric consistency (see Figure 3).

The natural discrete version of (i) is a closed, connected discrete surface in $X \times Y$.

Discrete version of (ii): As in the continuous setting we can project product triangles to the surfaces X and Y .

Definition 4. The projection $\pi_X : F \rightarrow \mathbb{Z}^{|F_X|}$ is defined by

$$\pi_X(f) := \begin{cases} e_a & \text{if } a = \begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix} \text{ is non-deg.} \\ (0, \dots, 0) & \text{else} \end{cases} \quad (13)$$

for each face $f = \begin{pmatrix} a_1, b_1 \\ a_2, b_2 \\ a_3, b_3 \end{pmatrix} \in F$. Here, e_a is the vector with 1 in the a -entry and 0 in all other entries.

We extend the projection π_X linearly to a map $\pi_X : \text{surf}(X \times Y) \rightarrow \mathbb{Z}^{|F_X|}$. The projection $\pi_Y : F \rightarrow \mathbb{Z}^{|F_Y|}$ and its linear extension $\pi_Y : \text{surf}(X \times Y) \rightarrow \mathbb{Z}^{|F_Y|}$ are defined similarly.

Let now Γ be a discrete surface in $X \times Y$. Then we say that the projections of Γ to X and Y are discrete diffeomorphisms if and only if

$$\begin{aligned} \pi_X(\Gamma) &= (1, \dots, 1) \in \mathbb{Z}^{|F_X|} \text{ and} \\ \pi_Y(\Gamma) &= (1, \dots, 1) \in \mathbb{Z}^{|F_Y|}. \end{aligned} \quad (14)$$

This gives a discrete version of (ii).

Note that in this definition we do not ask for injectivity on the vertices set. This is necessary for modelling discretely strong compressions. However, conditions (14) ensure a global bijectivity property which is sufficient in our context.

Discrete version of (iii): By definition, the set of surfaces in $X \times Y$ only contains surface patches which are consistently oriented. Therefore any surface in $\text{surf}(X \times Y)$ satisfies condition (iii).

Definition 5. Let $\Gamma \in \{0, 1\}^{|F|}$ be a discrete surface in $X \times Y$, represented by its indicator vector. Then Γ is a **discrete graph surface** in $X \times Y$ if

$$\begin{pmatrix} \partial \\ \pi_X \\ \pi_Y \end{pmatrix} \cdot \Gamma = \begin{pmatrix} \mathbf{0} \\ \mathbf{1} \\ \mathbf{1} \end{pmatrix}. \quad (15)$$

3.3. Discrete Surface Energy

Now we introduce a discrete energy on the set of product triangles in $X \times Y$. For the membrane energy in (2) we adopt the term proposed by Delingette [4]. Given two triangles $T_1, T_2 \subset \mathbb{R}^3$, Delingette computes the stretch energy $E_{\text{mem}}(T_1 \rightarrow T_2)$ necessary for deforming T_1 in T_2 . In our framework we associate with each product triangle

$(a, b) = \begin{pmatrix} a_1, b_1 \\ a_2, b_2 \\ a_3, b_3 \end{pmatrix} \in F$ the membrane cost

$$\begin{aligned} E_{\text{mem}}(a, b) &:= E_{\text{mem}} \left(\begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix} \rightarrow \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix} \right) + \\ &E_{\text{mem}} \left(\begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix} \rightarrow \begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix} \right). \end{aligned} \quad (16)$$

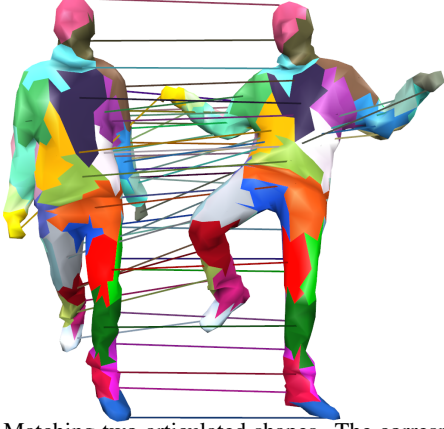


Figure 4. Matching two articulated shapes. The correspondences are visualized both by connecting lines and by patch coloring.

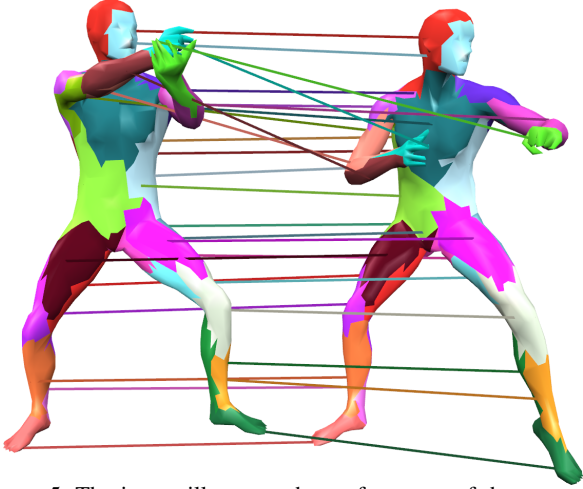


Figure 5. The image illustrates the performance of the proposed method on the SHREC 2011 [1] dataset.

For the bending term we proceed similarly associating with each product triangle (a, b) the cost

$$E_{\text{bend}}(a, b) = \int_a (H_X - H_Y)^2 + \int_b (H_Y - H_X)^2. \quad (17)$$

In practice we discretize the mean curvature following [17].

Next, we extend the energy linearly from discrete surface patches to discrete surfaces in $X \times Y$. Identify a discrete surface with its indicator vector $\Gamma \in \{0, 1\}^{|F|}$. Define the vector $E \in \mathbb{R}^{|F|}$ whose f -th entry is

$$E_f = E_{\text{mem}}(f) + E_{\text{bend}}(f). \quad (18)$$

Then the **discrete energy** of Γ is given by the vector product

$$E^t \cdot \Gamma. \quad (19)$$

4. Linear Programming Solution

In the previous section we have introduced a discrete notion of graph surfaces (15) and a discrete deformation en-

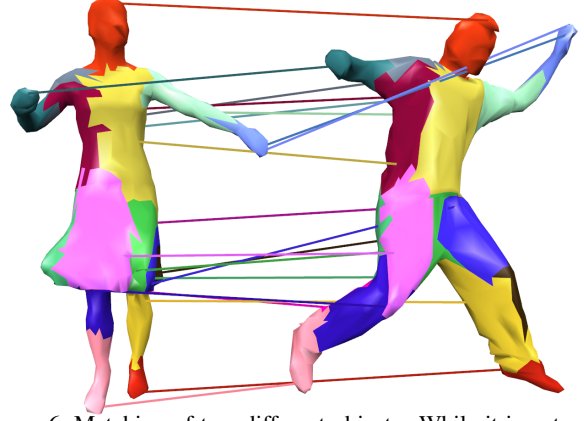


Figure 6. Matching of two different objects. While it is not well defined what a good matching between a skirt and trousers is, it is really remarkable how well the proposed algorithm finds a matching that apparently minimizes the deformation energy.

ergy (19) for such graph surfaces. This enables us to state the discrete version of (5) in the form of a binary linear program:

$$\begin{aligned} \min_{\Gamma \in \{0, 1\}^{|F|}} \quad & E^t \cdot \Gamma \\ \text{subject to} \quad & \begin{pmatrix} \partial \\ \pi_X \\ \pi_Y \end{pmatrix} \cdot \Gamma = \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix}. \end{aligned} \quad (20)$$

For solving (20), we relax the binary constraints to $\Gamma \in [0, 1]^{|F|}$. This relaxed version can be solved globally optimally in polynomial time. We employed an alternating direction method developed by Eckstein et al. [7]. This algorithm is parallelizable which allowed us an efficient implementation on the GPU.

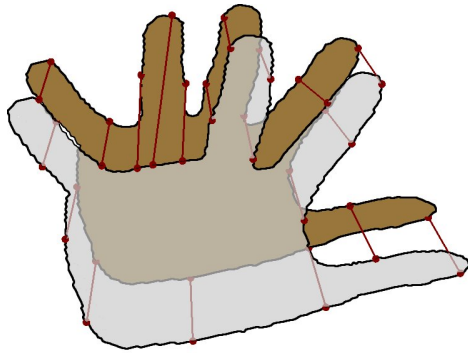
Since the constraint matrix of the relaxed problem is not totally unimodular, we are not guaranteed an integral solution. A simple thresholding scheme would destroy the geometric consistency of the solution. Therefore, for obtaining an integral solution we successively fix the variable with maximum value to 1.

5. Experimental Results

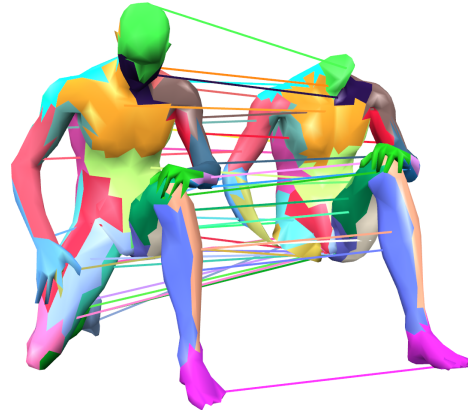
We have introduced a framework for computing geometrically consistent elastic matchings between 3D shapes using LP relaxation. We evaluated the proposed method on several shapes taken from the SHREC 2011 benchmark [1] and a dataset by Vlasic et al. [23].

5.1. Matching of Articulated Shapes

A common problem in shape matching is that the same shape may undergo substantial deformation and articulation. Nevertheless, one would like to reliably identify corresponding structures. Figures 4 and 8 show the matchings computed for models of different articulations. Although



Elastic matching of planar shapes [18]



Proposed elastic matching of 3D shapes

Figure 7. While the elastic matching of planar shapes can be solved in polynomial time as a minimal cyclic path on a torus [18], the framework developed in this paper allows to compute an elastic matching of 3D shapes via linear programming relaxation. In both cases, missing parts can be accounted for due to the elasticity.



Figure 8. The images show a matching between the leftmost and rightmost models taken from the SHREC 2011 benchmark [1] by linearly interpolating between the triangle correspondences. This transition illustrates the geometric consistency property of the proposed method: At any time during the interpolation the model is a closed triangle mesh.

the movement of arms and legs deform the shapes drastically the proposed method identifies the correct matchings. Since the proposed framework enforces geometric consistency matching errors occur only on a small spatial scale. In contrast to methods without spatial regularization strong outliers such as single points matched to the wrong leg do not arise.

5.2. Partial Matching

The ability of the proposed method to model stretching and shrinking also allows to match shapes where large parts of the geometry are missing. The right image of Figure 7 demonstrates this ability experimentally. The proposed algorithm matches the remaining parts of a human body missing a hand, a leg and the head to its original shape.

5.3. Quantitative Evaluation

We quantitatively evaluated the proposed method on 30 pairs of models from Vlasic et al. [23] by computing the mean geodesic error. One of the matchings is visualized in Figure 1. Computing each of the matchings took about 2

hours. The results were compared to matchings generated by the GMDS method of Bronstein et al. [2] using their code.

Given two meshes X, Y and the available ground truth correspondences (x_i, y_i) we defined the mean geodesic error of a matching $\varphi : X \rightarrow Y$ by $\frac{1}{N} \sum_i d(\varphi(x_i), y_i)$, where d is the normalized geodesic distance on the manifold of mesh Y .

The mean geodesic error produced by GMDS (using their code) was 0.079 while the proposed method had a mean geodesic error of 0.03.

Of course, this experiment does not pretend to be an exhaustive comparison against all methods in the literature. Nonetheless it shows, that the proposed method can compete with state-of-art matching algorithms in terms of accuracy while guaranteeing geometrically consistent solutions.

6. Conclusion

We proposed a new framework for finding a geometrically consistent matching of 3D shapes which minimizes an

elastic deformation energy. The approach is based on finding discrete minimal surfaces which set infinitesimal surface patches on both shapes into correspondence. In that sense the framework manages to generalize the planar shape matching to the more complex 3D shape matching. While the planar shape matching finds correspondences between infinitesimal line elements of two contours, the 3D shape matching sets infinitesimal surface elements in correspondence. We showed that a consistent discretization leads to an integer linear program. As a consequence, we can compute high-quality solutions to the matching problem which are independent of initialization by means of LP relaxation. Experimental results confirm that the proposed method generates reliable dense correspondences for a variety of articulated real-world shapes.

Acknowledgments

The 3D shape data in Figures 1, 4 and 6 is courtesy of Vlasic et al. [23]. The 3D shape data in Figures 5, 7 and 8 is taken from the SHREC 2011 benchmark [1].

References

- [1] E. Boyer, A. M. Bronstein, M. M. Bronstein, B. Bustos, T. Darom, R. Horaud, I. Hotz, Y. Keller, J. Keustermans, A. Kovnatsky, R. Litman, J. Reininghaus, I. Sipiran, D. Smeets, P. Suetens, D. Vandermeulen, A. Zaharescu, and V. Zobel. SHREC 2011: robust feature detection and description benchmark. *ArXiv e-prints*, February 2011. 6, 7, 8
- [2] A. Bronstein, M. Bronstein, and R. Kimmel. Efficient computation of isometry-invariant distances between surfaces. *SIAM J. Sci. Comput.*, 28(5):1812–1836 (electronic), 2006. 2, 7
- [3] P. Ciarlet. *An introduction to differential geometry with applications to elasticity*. Springer, Dordrecht, 2005. Reprinted from *J. Elasticity* 78/79 (2005), no. 1-3 [MR2196098]. 3
- [4] H. Delingette. Triangular springs for modeling nonlinear membranes. *IEEE Transactions on Visualisation and Computer Graphics*, 14(2), March/April 2008. 5
- [5] M. Desbrun, A. N. Hirani, M. Leok, and J. E. Marsden. Discrete exterior calculus, 2005. 4
- [6] M. Do Carmo. *Riemannian geometry*. Birkhauser, 1992. 3
- [7] J. Eckstein and D. Bertsekas. An alternating direction method for linear programming. Technical report, Harvard Business School, 1990. 6
- [8] M. Giaquinta, G. Modica, and J. Souček. *Cartesian currents in the calculus of variations. II*, volume 38 of *Ergebnisse der Mathematik und ihrer Grenzgebiete*. Springer-Verlag, Berlin, 1998. Variational integrals. 3
- [9] L. Grady. Minimal surfaces extend shortest path segmentation methods to 3d. *IEEE Trans. Pattern Anal. Mach. Intell.*, 32(2):321–334, 2010. 2
- [10] D. Hilbert. Über die Stetige Abbildung einer Linie auf ein Flächenstück. *Math. Ann.*, 38:459–460, 1891. 2
- [11] W. Koiter. On the nonlinear theory of thin elastic shells. I, II, III. *Nederl. Akad. Wetensch. Proc. Ser. B*, 69:1–17, 18–32, 33–54, 1966. 3
- [12] S. Kurtek, E. Klassen, Z. Ding, and A. Srivastava. A novel riemannian framework for shape analysis of 3d objects. In *CVPR*, pages 1625–1632, 2010. 2
- [13] Y. Lipman and I. Daubechies. Surface comparison with mass transportation, 2009. 2
- [14] N. Litke, M. Droske, M. Rumpf, and P. Schröder. An image processing approach to surface matching. In *Symposium on Geometry Processing*, pages 207–216, 2005. 2
- [15] L. Lui, T. Wong, P. Thompson, T. Chan, X. Gu, and S. Yau. Shape-based diffeomorphic registration on hippocampal surfaces using beltrami holomorphic flow. In *MICCAI (2)*, pages 323–330, 2010. 2
- [16] F. Mémoli and G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics*, 5(3):313–347, 2005. 2
- [17] M. Meyer, M. Desbrun, P. Schröder, and A. Barr. Discrete differential-geometry operators for triangulated 2-manifolds, 2002. 6
- [18] F. R. Schmidt, D. Farin, and D. Cremers. Fast matching of planar shapes in sub-cubic runtime. In *IEEE Int. Conf. on Computer Vision*, Rio de Janeiro, October 2007. 1, 3, 7
- [19] T. Schoenemann, F. Kahl, and D. Cremers. Curvature regularity for region-based image segmentation and inpainting: A linear programming relaxation. In *IEEE Int. Conf. on Computer Vision*, Kyoto, 2009. 2
- [20] J. M. Sullivan. *A Crystalline Approximation Theorem for Hypersurfaces*. PhD thesis, Princeton University, October 1990. 2
- [21] H. Tagare. Shape-based nonrigid correspondence with application to heart motion analysis. *IEEE Trans Med Imaging*, 18(7):570–579, 1999. 3
- [22] L. Torresani, V. Kolmogorov, and C. Rother. Feature correspondence via graph matching: Models and global optimization. In *ECCV (2)*, pages 596–609, 2008. 2
- [23] D. Vlasic, I. Baran, W. Matusik, and J. Popović. Articulated mesh animation from multi-view silhouettes. In *ACM SIGGRAPH 2008 papers*, SIGGRAPH '08, pages 97:1–97:9, New York, NY, USA, 2008. ACM. 6, 7, 8
- [24] Y. Wang, X. Gu, K. Hayashi, T. Chan, P. Thompson, and S. Yau. Brain surface parameterization using riemann surface structure. In *MICCAI (2)*, pages 657–665, 2005. 2
- [25] B. Wirth, L. Bar, M. Rumpf, and G. Sapiro. Geodesics in shape space via variational time discretization. In *EMM-CVPR'09*, volume 5681 of *LNCS*, pages 288–302, 2009. 2
- [26] L. Younes. *Shapes and diffeomorphisms*, volume 171 of *Applied Mathematical Sciences*. Springer-Verlag, Berlin, 2010. 2
- [27] Y. Zeng, C. Wang, Y. Wang, X. Gu, D. Samaras, and N. Paragios. Dense non-rigid surface registration using high-order graph matching. In *CVPR*, pages 382–389, 2010. 2