



Courbure discrète : théorie et applications

RENCONTRE ORGANISÉE PAR :
Laurent Najman and Pascal Romon

18-22 novembre 2013

Daniel Cremers, Emanuele Rodolà, and Thomas Windheuser

Relaxations for Minimizing Metric Distortion and Elastic Energies for 3D Shape Matching

Vol. 3, n° 1 (2013), p. 107-117.

<http://acirm.cedram.org/item?id=ACIRM_2013__3_1_107_0>

Centre international de rencontres mathématiques
U.M.S. 822 C.N.R.S./S.M.F.
Luminy (Marseille) FRANCE

cedram

*Texte mis en ligne dans le cadre du
Centre de diffusion des revues académiques de mathématiques
<http://www.cedram.org/>*

Relaxations for Minimizing Metric Distortion and Elastic Energies for 3D Shape Matching

Daniel CREMERS, Emanuele RODOLÀ, and Thomas WINDHEUSER

Abstract

We present two methods for non-rigid shape matching. Both methods formulate shape matching as an energy minimization problem, where the energy measures distortion of the metric defined on the shapes in one case, or directly describes the physical deformation relating the two shapes in the other case. The first method considers a parametrized relaxation of the widely adopted quadratic assignment problem (QAP) formulation for minimum distortion correspondence between deformable shapes. In order to control the accuracy/sparsity trade-off a weighting parameter is introduced to combine two existing relaxations, namely spectral and game-theoretic. This leads to an approach for deformable shape matching with *controllable sparsity*. The second method focuses on computing a geometrically consistent and spatially dense matching between two 3D shapes. Rather than mapping points to points it matches infinitesimal surface patches while preserving the geometric structures. In this spirit, matchings are considered 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, this leads to a minimal surface problem in this product space. The proposed discrete formulation describes the search space with linear constraints. Computationally, the approach results in a binary linear program whose relaxed version can be solved efficiently in a globally optimal manner.

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: In its most typical form, it concerns the problem of determining a map $f : X \rightarrow Y$ among two given shapes in such a way that their geometrical properties are preserved by the transformation. A particularly challenging instance of this problem occurs when the two shapes undergo general non-rigid deformations. As such, matching of deformable shapes has attracted the interest of researchers over the years and a wide variety of approaches have been proposed (see, e.g. [3] and references therein for a recent comparison).

A prominent approach to the matching problem from a metric perspective was introduced in [17], a concept that was explored further in [4] with the introduction of the GMDS framework, where the minimum distortion isometric embedding of one surface onto another is explicitly sought. A different view on the problem stems from the notion of uniformization space [14, 32]. Lipman and Funkhouser [14] proposed to model deviations from isometry by a transportation distance between corresponding points in a canonical domain (the complex plane); the result of this process is a

Text presented during the meeting “Discrete curvature: Theory and applications” organized by Laurent Najman and Pascal Romon. 18-22 novembre 2013, C.I.R.M. (Luminy).

Key words. shape matching, metric spaces, elastic deformation.

“fuzzy” correspondence matrix, whose values can be given the natural interpretation of confidence levels attributed to each match. Fuzzy schemes are typically adopted to relax the point-to-point mappings [18, 21]. Lipman and Daubechies [13] 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. In general, while methods based on uniformization theory are made attractive by the low dimensionality of the embedding domain, they do not behave well with different kinds of deformations (e.g., topological changes), and are subject to global inconsistencies in the final mapping.

In this work, we consider two different approaches to deformable shape matching. The two approaches share the common perspective of minimizing a distortion criterion, derived from the metric information which the shapes to be matched are endowed with. In one case (Section 3), following [23], we consider a notion of pairwise metric distortion that directly captures to what extent two shapes can be isometrically put in correspondence. Motivated by the observation that good correspondences often come at the price of high sparsity (in terms of number of matched points), whereas large cardinality tends to bring distorted matches into the correspondence, we attempt to control the accuracy/sparsity trade-off by introducing a weighting parameter on the combination of two effective relaxations [12, 21], which we relate to their regularizer counterparts from regression analysis. This leads us to the introduction of the *elastic net* penalty function [33] into shape matching problems. Differently, our second approach [30] takes a physically motivated view on the problem and minimizes a functional that encodes the physical deformation energy [15, 31] necessary to deform one shape into the other. The formulation we give in Section 4 is based on finding an optimal surface of codimension 2 in the product of the two shape surfaces. 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, our construction involves the boundary operator [27, 10, 25], and guarantees a geometrically consistent matching in the sense that the surfaces are mapped into one another in a continuous and orientation preserving manner.

2. ENERGY FUNCTIONALS FOR MEASURING THE MATCHING QUALITY

In this section we discuss the matching energies that have been used to find correspondences among shapes in [23] and [30] respectively.

2.1. Minimum metric distortion. We model shapes as compact Riemannian manifolds endowed with an intrinsic metric d . A point-to-point *correspondence* between two shapes X and Y is defined as a subset $C \subset X \times Y$ satisfying: 1) for every $x \in X$, there exists at least one $y \in Y$ such that $(x, y) \in C$, and vice versa, 2) for every $y \in Y$, there exists $x \in X$ such that $(x, y) \in C$. This relation can be alternatively formulated as a binary function $c : X \times Y \rightarrow \{0, 1\}$ satisfying the mapping constraints

$$(2.1) \quad \max_{x \in X} c(x, y) = \max_{y \in Y} c(x, y) = 1,$$

for every $y \in Y$ and $x \in X$. According to this definition, clearly not all correspondences give rise to meaningful matches among the two given shapes (consider, for instance, the full Cartesian product given by $c(x, y) = 1$ for all $(x, y) \in X \times Y$). A common requirement in this setting is that the correspondence should represent a bijective mapping, or more typically an *isometry* between the two surfaces. With this requirement in mind, in order to give a measure of *quality* to the correspondence we evaluate the distortion induced by the mapping as measured on the two shapes using the respective metrics d_X and d_Y . In particular, given two matches $(x, y), (x', y') \in C$, the absolute criterion

$$(2.2) \quad \epsilon(x, y, x', y') = |d_X(x, x') - d_Y(y, y')|$$

directly quantifies to what extent the estimated correspondence deviates from isometry. Following [18, 21], we first relax the correspondence from a discrete to a fuzzy notion by letting $c : X \times Y \rightarrow [0, 1]$, effectively setting off the problem from its combinatorial nature and bringing

it to a continuous optimization domain. Further, following a similar approach to the Gromov-Wasserstein [18] family of metrics, we obtain a *relaxed* notion of proximity between shapes:

$$(2.3) \quad D(X, Y) = \frac{1}{2} \min_C \sum_{(x,y),(x',y') \in C} e^p(x, y, x', y') c(x, y) c(x', y').$$

Note from this definition that we don't require the two shapes to have a measure defined over them (differently from [18, 21]). Establishing a minimum distortion correspondence between the two shapes amounts to finding a minimizer of the above distance. To this end, the problem can be easily recast as a relaxed quadratic assignment problem (QAP) [16],

$$(2.4) \quad \begin{aligned} \min_{\mathbf{C}} \quad & \text{vec}\{\mathbf{C}\}^T \mathbf{A} \text{vec}\{\mathbf{C}\} \\ \text{s.t.} \quad & \mathbf{C}\mathbf{1} = \mathbf{1}, \mathbf{C}^T \mathbf{1} = \mathbf{1}, \mathbf{C} \succeq 0, \end{aligned}$$

where $\text{vec}\{\mathbf{C}\}$ is the $|C|$ -dimensional column-stack vector representation of the correspondence matrix \mathbf{C} , \mathbf{A} is a non-negative symmetric cost matrix containing the pairwise distortion terms that appear in (2.3), $\mathbf{1}$ is a vector of $n = |C|$ ones, and \succeq denotes element-wise inequality. We emphasize that, although easier to solve, the relaxation provided above is still non-convex. Note that in the standard QAP, function c is taken to be a binary correspondence and matrix \mathbf{C} is thus required to be a permutation matrix. The QAP is a NP-hard problem due to the combinatorial complexity of this latter constraint.

2.2. Elastic deformation energies. A different approach to model a matching energy between shapes is to restrict the class of deformations that transform one shape into another to the set of diffeomorphisms. This gives us two benefits. First, the shapes do not get “cut” open during the matching transformation. Second, we can assign to each diffeomorphism an elastic energy that directly gives us a physical interpretation of the matching.

In the following, we assume that the two shapes $X, Y \subset \mathbb{R}^3$ are differentiable, oriented, closed surfaces. Diffeomorphisms $f : X \rightarrow Y$ are bijections for which both f and f^{-1} are differentiable. We formulate the shape matching problem as an optimization problem over the set of orientation preserving diffeomorphisms between X and Y ,

$$(2.5) \quad \inf_{f \in \text{Diff}^+(X, Y)} E(f) + E(f^{-1})$$

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 [5], which interprets 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:

$$(2.6) \quad E(f) = \underbrace{\int_X (\text{tr}_{g_X} \mathbf{E})^2 + \mu \text{tr}_{g_X} (\mathbf{E}^2)}_{E_{\text{mem}}} + \lambda \underbrace{\int_X (H_X(x) - H_Y(f(x)))^2}_{E_{\text{bend}}}$$

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 [8]), 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].

3. MINIMUM DISTORTION CORRESPONDENCE VIA ELASTIC NET REGULARIZATION

In this Section we present three different relaxations to the minimal metric distortion as formulated in problem (2.4). The three approaches act by relaxing the mapping constraints imposed on the correspondence function $c(x, y)$. Even though originating from distinct motivations, the first two methods share a convenient interpretation as partitioning problems in the space of potential assignments. In Section 3.3 we provide a different view on the problem, as presented in [23, 22], by using the language of regression analysis.

3.1. Spectral matching. Taking the point of view of graph clustering, [12] proposed the simplified problem

$$(3.1) \quad \min_{\mathbf{x}} \mathbf{x}^T \mathbf{A} \mathbf{x} \quad \text{s.t.} \|\mathbf{x}\|_2^2 = 1,$$

where $\mathbf{x} \equiv \text{vec}\{\mathbf{C}\} \in \mathbb{R}^n$ is the vector representation for the correspondence. Following Rayleigh’s quotient theorem, this modified QAP is minimized by the eigenvector \mathbf{x}^* corresponding to the minimum eigenvalue of \mathbf{A} . Note that mapping constraints are not imposed in (3.1). The authors follow a greedy algorithm to impose such constraints only after a solution has been obtained. The method has a tendency to produce matches for each point. Nevertheless, symmetries and structured noise in the data (indeed a characteristic of the non-rigid setting) may lead to unstable eigenvectors [12] and thus unreliable assignments.

A useful interpretation to this approach can be given as a relaxed two-way partitioning problem [1]. Consider the set of constraints taking the form $\mathbf{x}_i^2 = 1$ for $i = 1 \dots n$; these constraints restrict the values of \mathbf{x}_i to ± 1 , so the problem is equivalent to finding the partitioning (as “match” or “non-match”) on a set of n elements that minimizes the total cost $\mathbf{x}^T \mathbf{A} \mathbf{x}$. Here, the coefficients \mathbf{A}_{ij} can be interpreted as the cost of having elements i and j in the same partition. Clearly, the new constraints imply $\sum_{i=1}^n \mathbf{x}_i^2 = \|\mathbf{x}\|_2^2 = n$; since this actually allows the \mathbf{x}_i to take on any (small enough) real number, optimizing over this feasible set will yield a lower bound on the optimal value of the original partitioning problem.

3.2. Game-theoretic matching. Given the inherent difficulty to solve for a minimum distortion correspondence under general deformations, we recently proposed to shift the focus to the search of a group of matches having least distortion, *regardless* of its cardinality [21]. To achieve this, we proposed to optimize over the probability simplex

$$(3.2) \quad \|\mathbf{x}\|_1 = \mathbf{1}^T \mathbf{C} \mathbf{1} = 1, \quad \mathbf{x} \succeq 0.$$

In this formulation, the space of assignments is in a one-to-one correspondence with all possible probability distributions of a random variable, realizing as \mathbf{x} , modeling the concept of match. The main benefits of adopting such L^1 -type constraint for the matching problem arise from its convenient game-theoretical interpretation, leading to very efficient algorithms for (local) optimization and, most remarkably, in allowing the mapping constraints to be embedded directly into the cost matrix \mathbf{A} . Unfortunately, the strong locality and selectivity demonstrated by the game-theoretic approach is hardly desirable for matching problems.

Similarly to the L^2 case, the game-theoretic approach can be regarded as an attempt to solve a partitioning problem where the two partitions are represented by $\mathbf{x}_i = 0$ or 1 for $i = 1 \dots n$. This, in turn, corresponds to imposing a bound on the “counting” norm $\|\mathbf{x}\|_0$, which is relaxed here to the continuous sparsity-inducing counterpart $\sum_{i=1}^n |\mathbf{x}_i| = \|\mathbf{x}\|_1 = n$, with $\mathbf{x}_i \geq 0$ for all i .

3.3. Matching with the Elastic Net. In practical settings, the performance of the framework given in Section 2.1 directly depends on the definition of the metric distortion term ϵ . This is, in fact, a property shared by *any* method attempting to minimize (2.4). Ovsjanikov et al. [20] recently introduced the notion of shape condition number. According to this notion, the stability of the matching can be characterized as an intrinsic property of the shape itself, and is related to its intrinsic symmetries as well as the specific choice of a metric.

In order to incorporate a somewhat elusive notion of stability into the matching process, we propose to change the point of view by drawing an analogy between the correspondence problem and model-fitting. Our goal, in this context, is to determine a good approximation of the true relationship between the two shapes: we seek to *fit* or approximate the optimal correspondence \mathbf{x}^* as closely as possible, with deviation measured in the quadratic form $\mathbf{x}^T \mathbf{A} \mathbf{x}$. Problems of this kind are often studied with the tools of regression analysis [1]. Here the interest shifts from finding a best fit to analyzing the relationships among the several variables that build up the set of potential assignments $\{\mathbf{x}_i\}_{i=1 \dots n}$. These candidate matches act as predictors for the minimum distortion correspondence, and can be given the interpretation of explanatory variables which we observe, while we seek to find the combination that best describes the data in the minimal distortion sense. Since in general these variables hold a certain degree of correlation among them, it is of particular interest to attempt to determine whole groups of highly correlated predictors, as they will likely form consistent groups of matches in terms of the adopted measure of distortion.




Figure 3.1: Contour plots of the L^2 (circle), L^1 (diamond), and elastic net (in between) balls in \mathbb{R}^2 . In this example we set $\alpha = 0.6$. The strength of convexity varies with α .

In this view, spectral matching can be directly related to ridge regression, whereas the game-theoretic technique finds its equivalent in the lasso, the sparsity-inducing L^1 regularizer performing continuous shrinkage and automatic variable selection simultaneously [1, 33]; one major limitation of the lasso is its tendency to select only one variable from a group of variables among which the pairwise correlations are very high. In order to strike a balance between the two methods, we adopt a family of constraints known as elastic net [33]. This regularization technique shares with the lasso the ideal property of performing automatic variable selection, and most notably it is able to select entire groups of highly correlated variables. The elastic net criterion is defined as a convex combination of the lasso and ridge penalties:

$$(3.3) \quad (1 - \alpha)\|\mathbf{x}\|_1 + \alpha\|\mathbf{x}\|_2^2, \quad \alpha \in [0, 1].$$

This penalty function is singular at 0 and *strictly* convex for $\alpha > 0$, thus possessing the characteristics of both penalties (see Fig. 3.1). Strict convexity plays an important role as it guarantees the grouping effect in the extreme situation with identical predictors (that is, whenever the distortion between two matches is exactly 0), and provides a quantitative description of their degree of correlation otherwise. Let $\mathbf{x} \in \mathbb{R}^{|C|}$ be the vector representation of some correspondence $C \subset X \times Y$, we expect the elastic net-penalized solution to keep the difference $|\mathbf{x}_i - \mathbf{x}_j|$ small whenever the metric distortion $\epsilon(C_i, C_j)$ between the two matches is small. The trade-off between size of the correspondence and matching error is regulated by the convexity parameter α , which allows to fine tune the model complexity and balance the action of the penalty ranging from the highly selective pure lasso for $\alpha = 0$ to the more tolerant ridge behavior for $\alpha = 1$. This leads to the following family of relaxations for the QAP:

$$(3.4) \quad \begin{aligned} \min_{\mathbf{x}} \quad & \mathbf{x}^T \mathbf{A} \mathbf{x} \\ \text{s.t.} \quad & (1 - \alpha)\|\mathbf{x}\|_1 + \alpha\|\mathbf{x}\|_2^2 = 1, \quad \mathbf{x} \succeq 0, \end{aligned}$$

with $\alpha \in [0, 1]$. The family directly generalizes the spectral and game-theoretic techniques. Similarly to the spectral approach, this formulation does not guarantee the final solution to represent a bijective mapping, which can nevertheless be efficiently obtained *a posteriori* as in [12].

3.3.1. Optimization. We undertake a projected gradient approach [1] to determine a local optimum for problem (3.4). The optimization process is governed by the equations

$$(3.5) \quad \mathbf{x}^{(t+1)} = \Pi_{\alpha} \left(\mathbf{x}^{(t)} - \gamma^{(t)} \mathbf{A} \mathbf{x}^{(t)} \right),$$

where $\mathbf{A} \mathbf{x} = \frac{1}{2} \nabla \mathbf{x}^T \mathbf{A} \mathbf{x}$ is a descent direction for the objective, $\gamma > 0$ is the step length taken in that direction, and $\Pi_{\alpha} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a projection operator taking a solution back onto the feasible set. We initialize $\mathbf{x}^{(0)}$ to the barycenter of the elastic net boundary, i.e., for all $i = 1 \dots n$ we set \mathbf{x}_i to the positive solution of the quadratic equation $\alpha n x^2 + (1 - \alpha) n x - 1 = 0$.

While efficient methods for projecting onto the L^2 and L^1 balls have been proposed in literature [26], projection onto their convex combination is a more involved task. Computing the Euclidean projection $\Pi_{\alpha}(\mathbf{x}_0)$ onto the (positive) elastic net ball boundary amounts to solving the




Figure 3.2: Example of matchings obtained with the game-theoretic, elastic net and spectral techniques respectively. See Section 3.3.2 for details.

following problem:

$$(3.6) \quad \begin{aligned} \min_{\mathbf{x}} \quad & \|\mathbf{x} - \mathbf{x}_0\|_2^2 \\ \text{s.t.} \quad & (1 - \alpha)\mathbf{1}^T \mathbf{x} + \alpha \mathbf{x}^T \mathbf{x} = t, \quad \mathbf{x} \succeq 0, \end{aligned}$$

with $\alpha \in [0, 1]$. A detailed explanation of our approach on the computation of the unique minimum-distance projection Π_α in an efficient manner is given in [23]. Also note that, for practical purposes, we adopt a more efficient alternative to the standard projected gradient descent (3.5), namely its acceleration via vector extrapolation techniques [22].

3.3.2. Experimental results. We performed a wide range of experiments on the SHREC'10 standard dataset [3], which includes shapes undergoing several different types of deformation, e.g. quasi-isometric deformations, topological changes, displacement noise and changes in scale (we refer to [23] for a detailed numerical breakdown). Differently from most existing methods, the approach presented in this Section is quite general and not restricted to the quasi-isometric case. Indeed, invariance to different kinds of deformations is induced by the proper choice of the metrics employed in (2.2) (see [21] for an example). In order to make the computational task more tractable, only a limited number of samples are considered from one shape, and then potential matches are built with the 5 points from the other shape having similar curvature. Samples are generated via farthest point sampling (FPS) [17, 18] using the extrinsic Euclidean metric, a technique allowing to cover the whole surface in a sparse manner while retaining the metric information contained in the initial shape as best as possible. Note that only one of the two shapes is subsampled, while we keep all points in the other.

Fig. 3.2 presents an example in which the correct matches have a very small inlier ratio with respect to the set of candidates. In this matching scenario, our method provides a means to select only high-precision correspondences in a situation where there is huge ambiguity in most correspondences. In this example, the set of potential assignments is constructed by taking ~ 200 farthest points on one shape, and then building the whole Cartesian product with the *correct* corresponding points from the other shape, after 45% of them have been moved to random positions over the surface. This setup simulates a moderately challenging scenario in which only $\sim 50\%$ of the shape is matchable with low distortion, and the feasible set comprises all possible assignments between the two shapes. The game-theoretic (L^1) solution is highly selective and only assigns 3% of the shape samples accurately (left image); in contrast, the spectral (L^2) approach favors dense solutions and yields matches for 93% of the points with large error (right image). Elastic net matching (middle) allows to regulate the trade-off between size and distortion: the correspondence is made more dense, and 53% of the points are matched while keeping the error small. Here we set $\alpha = 0.85$.

4. MINIMIZING THE ELASTIC ENERGY VIA LINEAR PROGRAMMING RELAXATION

In this section we will discuss the approach presented in [30] that tries to solve the elastic energy problem

$$(4.1) \quad \inf_{f \in \text{Diff}^+(X, Y)} E(f) + E(f^{-1})$$

already introduced in Section 2.2. The approach puts the focus on three aspects:




Figure 4.1: The construction of product triangles that make up the discrete version of the product space $X \times Y$. **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 $(b_1, b_1, b_3)^T$.

- (1) Representation of the set of orientation-preserving diffeomorphisms $\text{Diff}^+(X, Y)$,
- (2) discretization of this set and the energy E , and
- (3) optimization of the discrete version of the energy.

The main idea underlying our representation is to look at subsets of the product space $X \times Y$. We will introduce constraints such that these subsets become graphs of diffeomorphisms. We will show further how we can discretize the product space, the constraints and the energy. Interestingly, the constraints and the energy are linear in the variables that span the discretized version of $X \times Y$. The resulting optimization problem is thus an *integer linear program* (ILP).

While we cannot find the global optimum of this optimization problem we can allow non-integer solutions and transform the ILP into a *linear program* (LP). The global optimum of the linear program can be computed in polynomial time and is a lower bound of the original optimization problem.

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

$$(4.2) \quad \Gamma = \{(x, f(x)) \mid x \in X\} \subset X \times Y.$$

The set 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.

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

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

$$(4.3) \quad E(f) = \tilde{E}(\Gamma)$$

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 (4.1) can be phrased as an optimization problem over the set of subsets of $X \times Y$, which then reads

$$(4.4) \quad \begin{aligned} & \inf \quad \tilde{E}(\Gamma) \\ & \text{subject to } \Gamma \subset X \times Y \text{ is a graph surface} \end{aligned}$$

We remark that the idea of casting optimal diffeomorphism problems as minimal surface problems has been applied previously in the theory of nonlinear elasticity [9]. In the setup of shape matching, it is related to the approach that Tagare [28] proposed for the matching of 2D shapes. It was reformulated as an orientation preserving diffeomorphism approach in [24].




Figure 4.2: The discrete version of condition (i) includes the closeness condition ensuring that neighboring triangles on X are matched with neighboring triangles on Y . **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 closeness condition is granted by the boundary operator evaluated on the product edges $((a_2, b_1), (a_3, b_3))^T$ and $((a_3, b_3), (a_1, b_1))^T$.

4.2. The discrete setting. We develop now a discrete counterpart of the notion of graph surfaces in $X \times Y$ and the continuous elastic matching energy by assuming that the surfaces X, Y are given as triangulated meshes.

4.2.1. 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 $\begin{pmatrix} a_1 \\ a_2 \end{pmatrix} \in E_X$ or $\begin{pmatrix} a_2 \\ a_1 \end{pmatrix} = -\begin{pmatrix} a_1 \\ a_2 \end{pmatrix} \in E_X$. We extend the set of edges by *degenerate edges* $\bar{E}_X = E_X \cup \left\{ \begin{pmatrix} a \\ a \end{pmatrix} \mid a \in V_X \right\}$. 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 $\begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix} = \begin{pmatrix} a_2 \\ a_3 \\ a_1 \end{pmatrix} = \begin{pmatrix} a_3 \\ a_1 \\ a_2 \end{pmatrix} \in F_X$. Similarly, we extend the set of triangles by *degenerate triangles* $\bar{F}_X = F_X \cup \left\{ \begin{pmatrix} a_1 \\ a_2 \\ a_2 \end{pmatrix} \mid a_1, a_2 \in V_X, \pm \begin{pmatrix} a_1 \\ a_2 \end{pmatrix} \in \bar{E}_X \right\}$. Notice that degenerate triangles can consist of only one or two vertices. The existence of these degenerate triangles will allow stretching or compression of parts of the surface.

Next, we introduce product triangles for two triangulated 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

$$(4.5) \quad F := \left\{ \begin{array}{l} \left(\begin{array}{l} (a_1, b_1) \\ (a_2, b_2) \\ (a_3, b_3) \end{array} \right) \left| \begin{array}{l} f_1 = \begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix} \in \bar{F}_X, \\ f_2 = \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix} \in \bar{F}_Y, \\ f_1 \text{ or } f_2 \text{ non-degenerate} \end{array} \right. \end{array} \right\}$$

We will call the triple (V, E, F) the *product graph* as the discrete counterpart of the product space $X \times Y$. The product triangles in F are the basic pieces which are later glued to discrete graph surfaces. For shape matching, a product triangle $((a_1, b_1), (a_2, b_2), (a_3, b_3)) \in F$ is interpreted as setting vertex $a_i \in V_X$ in correspondence with vertex $b_i \in V_Y$ (see Figure 4.1).

4.2.2. Discrete surfaces. Following Proposition 1 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 the discrete counterpart of a diffeomorphic 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 [7] 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 $e \in E$. We then define for two vertices $v_1, v_2 \in V$ a vector $O \binom{v_1}{v_2} \in \mathbb{Z}^{|E|}$ whose e -th entry is given by

$$(4.6) \quad O \binom{v_1}{v_2}_e = \begin{cases} 1 & \text{if } e = \binom{v_1}{v_2} \\ -1 & \text{if } e = \binom{v_2}{v_1} \\ 0 & \text{else.} \end{cases}$$

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 defined by

$$(4.7) \quad \partial \begin{pmatrix} a_1, b_1 \\ a_2, b_2 \\ a_3, b_3 \end{pmatrix} := O \begin{pmatrix} a_1, b_1 \\ a_2, b_2 \end{pmatrix} + O \begin{pmatrix} a_2, b_2 \\ a_3, b_3 \end{pmatrix} + O \begin{pmatrix} a_3, b_3 \\ a_1, b_1 \end{pmatrix},$$

where the $a_i \in V_X$ and $b_i \in V_Y$ form triangles on X resp. on Y and $\binom{a_i, b_i}{a_j, b_j}$ 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 $\partial : \text{surf}(X \times Y) \rightarrow \mathbb{Z}^{|E|}$. 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 a discrete notion of continuity (see Figure 4.2). 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 by defining $\pi_X : F \rightarrow \mathbb{Z}^{|F_X|}$ as

$$(4.8) \quad \pi_X(f) := \begin{cases} e_a & \text{if } a = \begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix} \text{ is non-deg.} \\ (0, \dots, 0) & \text{else} \end{cases}$$

for each face $f = ((a_1, b_1), (a_2, b_2), (a_3, b_3)) \in F$. Here, e_a is the vector with 1 in the a -entry and 0 in all other entries. We extend the projection linearly to $\pi_X : \text{surf}(X \times Y) \rightarrow \mathbb{Z}^{|F_X|}$.

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 $\pi_X(\Gamma) = (1, \dots, 1) \in \mathbb{Z}^{|F_X|}$ and $\pi_Y(\Gamma) = (1, \dots, 1) \in \mathbb{Z}^{|F_Y|}$. 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, they 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).

4.2.3. Discrete surface energy. Now we introduce a discrete energy on the set of product triangles in $X \times Y$. For the membrane energy in (2.6) we adopt the term proposed by Delingette [6]. 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 make the energy symmetric by associating with each product triangle $(a, b) \in F$ the membrane cost $E_{\text{mem}}(a, b) := E_{\text{mem}}(a \rightarrow b) + E_{\text{mem}}(b \rightarrow a)$. For the bending term we proceed similarly associating with each product triangle (a, b) the cost $E_{\text{bend}}(a, b) = \int_a (H_X - H_Y)^2 + \int_b (H_Y - H_X)^2$. In practice we discretize the mean curvature following [19].

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 $\Gamma \in \{0, 1\}^{|F|}$, and define the vector $E \in \mathbb{R}^{|F|}$ whose f -th entry is $E_f = E_{\text{mem}}(f) + E_{\text{bend}}(f)$. Then the *discrete energy* of Γ is given by the vector product $E^T \Gamma$.




Figure 4.3: Matchings between 3D shapes obtained by minimizing the relaxation of energy (4.9). Since the energy can express elastic deformations such as stretching and shrinking the proposed method can find transformations that are highly non-rigid and non-isometric. The 3D shape data is from Vlasic et al. [29] (left) and the SHREC 2011 benchmark [2] (right).

4.2.4. *Optimizing the discrete energy.* The notion of discrete graph surfaces and the discrete surface energy introduced in Sections 4.2.2 and 4.2.3 can be combined with the discrete version of optimization problem (4.4) in the form of a binary linear program:

$$(4.9) \quad \begin{aligned} \min_{\Gamma \in \{0,1\}^{|F|}} \quad & E^T \Gamma \\ \text{subject to} \quad & \begin{pmatrix} \partial \\ \pi_X \\ \pi_Y \end{pmatrix} \Gamma = \begin{pmatrix} \mathbf{0} \\ \mathbf{1} \\ \mathbf{1} \end{pmatrix}. \end{aligned}$$

Similarly to what we did in (2.3), in order to solve (4.9) we relax the binary constraints to $\Gamma \in [0,1]^{|F|}$. For this relaxed version the global optimum can be computed in polynomial time. 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. Typical matching results are shown in Figure 4.3. For a more detailed experimental evaluation we refer to [30].

5. CONCLUSIONS

In this paper we discussed two approaches to non-rigid shape matching by optimizing a distortion criterion. While for both approaches the distortion criterion is based on information derived from the metric of the shape, it is possible to express quite different notions of similarity, i.e. a geometrical Gromov-Wasserstein distance and a physical thin-shell energy. By following different algorithmic strategies for both notions of similarity, we showed that it is possible to find good matchings minimizing the distortion energies between the shapes.

ACKNOWLEDGMENTS

We thankfully acknowledge Matthias Vestner for helpful discussions.

REFERENCES

- [1] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge Univ. Press, New York, USA, 2004.
- [2] 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.
- [3] Alex Bronstein, Michael Bronstein, Umberto Castellani, et al. Shrec 2010: Robust correspondence benchmark. In *Eurographics Workshop on 3D Object Retrieval*, 2010.
- [4] Alex Bronstein, Michael Bronstein, and Ron Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proc. National Academy of Science (PNAS)*, 103(5):1168–1172, 2006.
- [5] 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].

- [6] H. Delingette. Triangular springs for modeling nonlinear membranes. *IEEE Transactions on Visualisation and Computer Graphics*, 14(2), March/April 2008.
- [7] M. Desbrun, A. N. Hirani, M. Leok, and J. E. Marsden. Discrete exterior calculus, 2005.
- [8] M.P. Do Carmo. *Riemannian geometry*. Birkhauser, 1992.
- [9] 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.
- [10] Leo Grady. Minimal surfaces extend shortest path segmentation methods to 3d. *IEEE Trans. Pattern Anal. Mach. Intell.*, 32(2):321–334, 2010.
- [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.
- [12] Marius Leordeanu and Martial Hebert. A spectral technique for correspondence problems using pairwise constraints. In *Proc. CVPR*, volume 2, pages 1482–1489, 2005.
- [13] Y. Lipman and I. Daubechies. Surface comparison with mass transportation, 2009.
- [14] Yaron Lipman and Thomas Funkhouser. Mobius voting for surface correspondence. *ACM Trans. on Graphics*, 28(3), 2009.
- [15] 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.
- [16] Eliane Maria Loiola, Nair Maria Maia de Abreu, Paulo Oswaldo Boaventura-Netto, Peter Hahn, and Tania Querido. A survey for the quadratic assignment problem. *European Journal of Operational Research*, 176(2):657–690, 2007.
- [17] F. Mémoli and G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Found. Comput. Math.*, 5:313–346, 2005.
- [18] Facundo Mémoli. Gromov-Wasserstein distances and the metric approach to object matching. *Found. Comput. Math.*, 11:417–487, 2011.
- [19] M. Meyer, M. Desbrun, P Schröder, and A. Barr. Discrete differential-geometry operators for triangulated 2-manifolds, 2002.
- [20] Maks Ovsjanikov, Qi-Xing Huang, and Leonidas J. Guibas. A condition number for non-rigid shape matching. *Comput. Graph. Forum*, pages 1503–1512, 2011.
- [21] Emanuele Rodolà, Alex Bronstein, Andrea Albarelli, Filippo Bergamasco, and Andrea Torsello. A game-theoretic approach to deformable shape matching. In *Proc. CVPR*, 2012.
- [22] Emanuele Rodolà, Tatsuya Harada, Yasuo Kuniyoshi, and Daniel Cremers. Efficient shape matching using vector extrapolation. In *Proc. BMVC*, 2013.
- [23] Emanuele Rodolà, Andrea Torsello, Tatsuya Harada, Yasuo Kuniyoshi, and Daniel Cremers. Elastic net constraints for shape matching. In *Proc. ICCV*, Sydney, Australia, December 2013.
- [24] F. R. Schmidt, Dirk 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.
- [25] 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.
- [26] Shai Shalev-Shwartz and Yoram Singer. Efficient learning of label ranking by soft projections onto polyhedra. *J. Mach. Learn. Res.*, 7:1567–1599, December 2006.
- [27] John M. Sullivan. *A Crystalline Approximation Theorem for Hypersurfaces*. PhD thesis, Princeton University, October 1990.
- [28] Hermant Tagare. Shape-based nonrigid correspondence with application to heart motion analysis. *IEEE Trans Med Imaging*, 18(7):570–579, 1999.
- [29] Daniel Vlastic, Ilya Baran, Wojciech Matusik, and Jovan 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.
- [30] Thomas Windheuser, Ulrich Schlickewei, Frank R Schmidt, and Daniel Cremers. Geometrically consistent elastic matching of 3d shapes: A linear programming solution. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 2134–2141. IEEE, 2011.
- [31] B. Wirth, L. Bar, M. Rumpf, and G. Sapiro. Geodesics in shape space via variational time discretization. In *EMMCVPR'09*, volume 5681 of *LNCS*, pages 288–302, 2009.
- [32] Yun Zeng, Chaohui Wang, Yang Wang, Xianfeng Gu, Dimitris Samaras, and Nikos Paragios. Dense non-rigid surface registration using high-order graph matching. In *Proc. CVPR*, pages 382–389, 2010.
- [33] Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society, Series B*, 67:301–320, 2005.