Regularized Partial Matching of Rigid Shapes

Alexander M. Bronstein and Michael M. Bronstein

Technion – Israel Institute of Technology, Haifa 32000, Israel
{bron,mbron}@cs.technion.ac.il

Abstract. Matching of rigid shapes is an important problem in numerous applications across the boundary of computer vision, pattern recognition and computer graphics communities. A particularly challenging setting of this problem is partial matching, where the two shapes are dissimilar in general, but have significant similar parts. In this paper, we show a rigorous approach allowing to find matching parts of rigid shapes with controllable size and regularity. The regularity term we use is similar to the spirit of the Mumford-Shah functional, extended to non-Euclidean spaces. Numerical experiments show that the regularized partial matching produces better results compared to the non-regularized one.

1 Introduction

Shape matching is one of the cornerstone problems arising in many applications across the boundaries of computer vision, shape recognition and computer graphics communities. By saying *matching*, we actually refer here to two separate problems: *similarity* and *correspondence*. The first problem consists of computing a "distance" between the shapes and is usually dealt with in computer vision and pattern recognition applications. The purpose of the second is finding the matching points between the shapes; it arises in computer graphics community in shape synthesis applications. Usually, the two problems are inter-related and one can be solved as a byproduct of the other.

Many matching algorithms exist in the literature depending on the nature of shapes and their representation [17]. A more complicated setting of shape matching is encountered when the shapes have similar parts but are dissimilar as a whole. For example, the shapes of a centaur and a human fall into this category [9]. A semantically correct matching between such shapes would be the upper part of the body, which both the centaur and the human share. We call this type of matching *partial*. Partial matching plays a crucial role in many practical problems, in which the data to be matched are not available entirely due to acquisition imperfections. For example, if one wishes to match two instances of the same shape acquired using a three-dimensional scanner, the shapes may be different due to scanning artifacts (noise, holes, etc.) It is often important to have control over the selected parts, e.g. controlling their size and regularity.

In computer graphics applications, rigid shape matching is usually performed using different flavors of the classical *iterative closest point* (ICP) algorithm [6,2,12,8]. ICP algorithms try to optimally align the shapes by finding a rigid transformation minimizing a surface-to-surface distance between them. Partial matching can be performed using ICP with the rejection or down-weighting of points with a "bad" correspondence. This approach allows to find best matching parts of shapes, but does not allow direct

D. Forsyth, P. Torr, and A. Zisserman (Eds.): ECCV 2008, Part II, LNCS 5303, pp. 143–154, 2008.

control of their size and regularity. Another common family of approaches is based on finding a set of features on the shapes, followed by computation of invariant local descriptor for each feature. The reader is referred to [17] for a comprehensive review of these methods. A recent study on using local descriptor-based approaches for partial shape similarity is presented in [7].

In the computer vision community, the problem of partial matching of objects is sometimes approached using the "recognition by parts" idea [15]: segmenting the shape into significant parts and matching pairs of parts as whole shapes [10,16,3]. The main difficulty of this approach is that the definition of "significant parts" is semantic rather than geometric, and therefore, automatically finding such parts is not a well-defined problem [1]. Trying to avoid this problem, Latecki *et al.* [11] proposed a partial similarity criterion that removed the arbitrarity of shape partition. A meaningful part is defined as the most similar common part of two shapes, and is practically found by simplifying the shapes until they look the most similar. Bronstein *et al.* [4] extended this idea, proposing a Paretian framework for the computation of partial matching of *nonrigid* shapes. Their approach consists of solving a multicriterion optimization problem trying to simultaneously maximize the similarity and the size of the parts and allows control the size of the selected parts. The main disadvantage of this approach is that it ignores the shape of the selected parts, which theoretically can be irregular.

In this paper, we show that the method of [4] can also be used for partial matching of rigid shapes and improve it by taking into account the regularity of parts besides their size. The regularity term proposed here is similar to the spirit of the Mumford-Shah [13] and Chan-Vese [5] functionals. The rest of the paper is organized as follows. In Section 2, we present the problem of partial rigid shape matching and iterative closest point (ICP) algorithms with rejection of points attempting to solve this problem. We point out the main drawback of this method: the part size and regularity are not directly controllable. In Section 3, we present our approach for regularized partial shape matching with controllable part size and regularity and in Section 4 describe its numerical computation. Section 5 shows experimental results. Finally, Section 6 concludes the paper.

2 Problem Definition

Let X and Y be two rigid shapes, considered here as a subsets of the Euclidean space $\mathbb{R}^3$. The shapes are said to be *congruent* if applying a *rigid motion* to one of the shape we can obtain the other, or in other words, $X = RY + t$ (here R is a rotation matrix and t is a translation vector). Using congruence as an equivalence relation, we say that two rigid shapes are equivalent if they are congruent.

In practice, shapes rarely happen to be precisely congruent. We can still use the degree of incongruence of two shapes in order to quantify their dissimilarity. It is common to use the *Hausdorff distance*

$$d_{\mathrm{H}}(X,Y) = \max\left\{ \max_{x\in X} \min_{y\in Y} \|x-y\| \,,\, \max_{y\in Y} \min_{x\in X} \|x-y\| \right\} \tag{1}$$

to measure the distance between to subsets of $\mathbb{R}^3$. Trying to minimize d_{H} over all the possible rigid motions, we obtain a criterion of rigid shape dissimilarity,

$$d(X,Y) = \min_{R,t} d_{\mathrm{H}}(X, RY + t) \tag{2}$$

X and Y are congruent if and only if $d(X,Y) = 0$.

Defining $y^* : X \to Y$ to be the function mapping a point in x on X to the closest point $y^*(x) = \arg\min_{y \in Y} \|x-y\|^2$ on Y, and, analogously, $x^*(y) = \arg\min_{x \in X} \|x-y\|^2$, to be the closest point on X, the Hausdorff distance can be rewritten as

$$d_{\mathrm{H}}(X,Y) = \max\left\{ \max_{x \in X} \|x - y^*(x)\|, \max_{y \in Y} \|y - x^*(y)\| \right\}. \tag{3}$$

Using this formulation, the distance in (2) can be computed using an iterative two-stage process: first, find the correspondence between X and Y. Second, the optimal rigid motion parameters R, t are found to minimize the distance (3) between the closest points on X and Y. The process is repeated until convergence. This class of method is commonly known as *iterative closest point* (ICP) algorithms [6,2]. The outcome of these algorithms is two-fold: the correspondences $x^*(Y)$ and $y^*(X)$ and the dissimilarity $d(X,Y)$ between the shapes X and Y. In practice, the L_∞ distance (3) is usually replaced by an L_2 version,

$$d_{\mathrm{H}}(X,Y) = \int_X \|x - y^*(x)\|^2 dx + \int_Y \|y - x^*(y)\|^2 dy, \tag{4}$$

or its non-symmetric variant, involving only the first term.

2.1 Partially Similar Shapes

In some cases, it may happen that while the shapes X and Y are dissimilar as whole, their parts are similar. Using straightforwardly rigid matching methods will produce meaningless results, failing to capture the partial similarity of shapes. In order to find a partial matching between X and Y, we need to determine similar parts $X' \subseteq X$ and $Y' \subseteq Y$. A common method used for this purpose is a modification of the ICP algorithm with rejection or down-weighting of dissimilar point pairs. Using this approach, the Hausdorff distance in the ICP algorithm is replaced for example by

$$d_{\mathrm{H}}(X,Y) = \int_X \|x - y^*(x)\|^2 w(x) dx, \tag{5}$$

where $w(x)$ is a weighting function rejecting points with bad correspondence. The "quality" of correspondence can be determined by the distance between x and $y^*(x)$ or the mismatch between the normals to X and Y at the points x and $y^*(x)$, or a combination of both. For example, $w(x)$ can be defined as

$$w(x) = \begin{cases} 1 : \langle n_X(x), n_Y(y^*(x)) \rangle \geq \tau_{\mathrm{n}} \text{ and } \|x - y^*(x)\| \leq \tau_{\mathrm{d}} \\ 0 : \text{else} \end{cases} \tag{6}$$

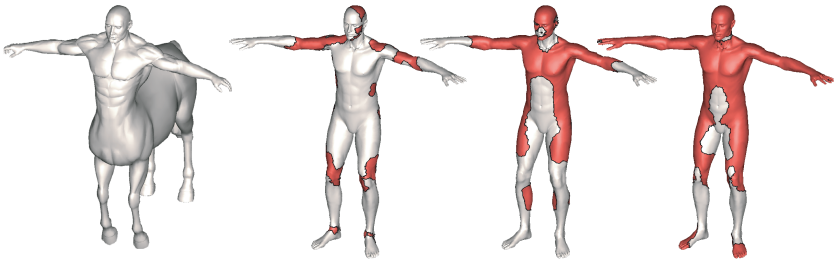

Fig. 1. Matching between a centaur (left) and a man (three rightmost shapes) using ICP with different degree of rejection of points, corresponding approximately to 25%, 50%, and 75% of the shape area (left-to-right). Note that the obtained parts suffer from irregularity.

where τ_{n} and τ_{d} are some fixed thresholds. Points for which $w = 0$ will be rejected, i.e., will not contribute to $d_{\mathrm{H}}(X, Y)$. The matching parts are defined as $X' = \{x \in X : w(x) = 1\}$ and $Y' = y^*(X')$.

By changing the rejection thresholds τ_{n} and τ_{d}, it is possible to control the size (area) of the parts (Figure 1). However, this control is indirect, as there is no explicit relation between the surface-to-surface distance and the area of the parts – the size of the parts is known only *post factum*. Another drawback of this method is evident from Figure 1: the matching parts are irregular and contain multiple disconnected components and holes. The influence of the rejection thresholds on the part size and regularity is indirect and depends on the specific shapes, and thus is not explicitly controllable.

3 Regularized Partial Matching of Shapes

The main contribution of this paper is a framework for partial matching of shapes with controllable size and regularity of parts. We call this approach *regularized partial matching*. Given two partially similar shapes X and Y, our goal is to select two parts $X' \subseteq X$ and $Y' \subseteq Y$, such that they are as regular, similar and large as possible. It is clear that the three criteria are conflicting. For example, increasing the size of the parts makes them less similar, and in order to make the parts regular we may sacrifice both size and similarity.

More formally, we have a multicriterion optimization problem with three objectives: dissimilarity $d_{\mathrm{H}}(X', Y')$, *partiality* $p(X')+p(Y') = \mathrm{area}(X)-\mathrm{area}(X')+\mathrm{area}(Y)-\mathrm{area}(Y')$ and *irregularity* $r(X') + r(Y')$, which we wish to simultaneously minimize over all the possible parts,

$$\min_{X' \subseteq X, Y' \subseteq Y} (d(X', Y'), p(X') + p(Y'), r(X') + r(Y')). \tag{7}$$

The solution of the multicriterion optimization problem (7) is the set of parts (X^*, Y^*) achieving an optimal tradeoff between the criteria, in the sense that there exists no other pair of parts (X', Y') with $d_{\mathrm{H}}(X', Y') < d(X^*, Y^*), p(X')+p(Y') < p(X^*)+p(Y^*)$ and $r(X')+r(Y') < r(X^*)+r(Y^*)$ holding simultaneously. Such a solution is called

Pareto optimal and is not unique. One of the ways to solve the above problem is by fixing some partiality p_0 and minimizing

$$\min_{X' \subseteq X, Y' \subseteq Y} d(X', Y') + \mu(r(X') + r(Y')) \text{ s.t. } p(X') + p(Y') \leq p_0, \tag{8}$$

where λ is a non-negative coefficient, determining the relative importance of the similarity and regularity of the parts.

As the criterion of part irregularity, the simplest choice is the length of the part boundary $\partial X'$,

$$r(X') = \int_{\partial X'} d\ell. \tag{9}$$

Using this definition, the minimization problem of the part irregularity $r(X')$ subject to fixed partiality $p(X') = p_0$ can be regarded as an *isoperimetric problem*. In case of two-dimensional shapes (subsets of the Euclidean plane) the minimum is achieved by a circle, which in many applications is undoubtfully the most regular shape. Unfortunately, there exists no known extension of this result to curved surfaces, on which, in general, we may find two parts with the same area and boundary length having an arbitrarily large number of disconnected components. In the case when the topological regularity of the part is important, $r(X')$ can be defined as the *genus* of X', which according to the Gauss-Bonnet theorem can be expressed as

$$r(X') = 1 - \frac{1}{2}\int_{X'} K(x)dx - \frac{1}{2}\int_{\partial X'} \kappa_g(x)d\ell, \tag{10}$$

where K and κ_g are the Gaussian and geodesic curvatures, respectively. Such a definition will penalize parts having holes or multiple disconnected components (and, thus, larger genus), and provide the desired topological regularity. In general, different irregularity criteria can be defined to suit the specific application needs. In the sequel, we will stick to the boundary length irregularity (9) for simplicity.

3.1 Fuzzy Formulation

The solution of (8) involves minimization over the set of all pairs of parts of X and Y, which can be thought of as minimization over all pairs of binary *membership functions* $u : X \to \{0, 1\}$ and $v : Y \to \{0, 1\}$, specifying for each point in X and Y whether it belongs to the part or not. Such a discrete minimization problem is clearly computationally intractable. As a remedy, we replace the binary membership functions u and v by a *fuzzy* approximation $u : X \to [0, 1]$ and $v : Y \to [0, 1]$, bringing the problem back to a tractable continuous formulation. We will denote by $\tilde{X} = (X, u)$ and $\tilde{Y} = (Y, v)$ the *fuzzy parts* of X and Y, respectively. In this formulation, the dissimilarity between two parts is given by

$$d(\tilde{X}, \tilde{Y}) = \min_{R,t} d_{\mathrm{H}}(\tilde{X}, R\tilde{Y} + t), \tag{11}$$

where $R\tilde{Y} + t = (RY + t, v)$, and

$$d_{\mathrm{H}}(\tilde{X}, \tilde{Y}) = \int_X \|x - y^*(x)\|^2 u(x)dx + \int_Y \|y - x^*(y)\|^2 v(y)dy. \tag{12}$$

The sets of the closest points $x^*(Y)$ and $y^*(X)$ are defined as before, and are functions of X and Y. Denoting by $e_{X,Y}(x) = \|x - y^*(x)\|^2$ and $e_{Y,X}(y) = \|y - x^*(y)\|^2$ the local measures of surface misalignment, we can formulate the dissimilarity as

$$d(\tilde{X}, \tilde{Y}) = \min_{R,t} \int_X e_{X,RY+t}(x)u(x)dx + \int_Y e_{RY+t,X}(y)v(y)dy. \tag{13}$$

Here, for the sake of simplicity, we will stick to these basic definitions of surface misalignment. In practice, other measures involving, for example, the misalignment of the normals of X and Y can be used as well.

To conclude our discussion on fuzzy approximation, we need to "fuzzify" the partiality and the irregularity terms. The partiality of a fuzzy part $\tilde{X}$ straightforwardly becomes

$$p(\tilde{X}) = \int_X (1 - u(x))dx. \tag{14}$$

The irregularity term is slightly more elaborate, since in the fuzzy part there is no more boundary in the strict sense. However, adopting the Mumford-Shah spirit, we can replace integration along the boundary by integration of the band in which the membership function changes from small to large values [13,5],

$$r(\tilde{X}) = \int_X h(u(x)) \|\nabla_X u(x)\| dx \tag{15}$$

where $h(t) \approx \delta(t - 0.5)$ is an approximation of the Dirac delta function, and $\nabla_X u(x)$ is the intrinsic gradient of u at the point x. The quantity $\|\nabla_X u(x)\|$ can be thought of as the length of the extrinsic gradient vector $\nabla_{\mathbb{R}^3} u$ projected on the tangent space of X at a point x.

3.2 Alternating Minimization Algorithm

Substituting the fuzzy approximation of dissimilarity, partiality and regularity into the minimization problem (8), we obtain

$$\begin{aligned} \min_{R,t,u,v} & \int_X (e_{X,RY+t}(x)u(x) + \mu h(u(x)) \|\nabla_X u(x)\|)dx + \\ & \int_Y (e_{RY+t,X}(y)v(y) + \mu h(v(y)) \|\nabla_Y v(y)\|)dy \\ \text{s.t.} & \int_X u(x)dx + \int_Y v(y)dy \geq (1 - p_0)(\text{area}(X) + \text{area}(Y)). \end{aligned} \tag{16}$$

Here, for convenience, we normalized the partiality parameter p_0 by the total area of X and Y. By fixing the membership functions u and v, the minimization over R and t becomes

$$\min_{R,t} \int_X e_{X,RY+t}(x)u(x)dx + \int_Y e_{RY+t,X}(y)v(y)dy, \tag{17}$$

which can be solved using a weighted ICP algorithm. On the other hand, fixing R and t yields the linearly constrained nonlinear minimization problem

$$\begin{aligned}\min_{u,v}\ & \int_X (e_{X,Y}(x)u(x) + \mu h(u(x))\,\|\nabla_X u(x)\|)dx + \\ & \int_Y (e_{Y,X}(y)v(y) + \mu h(v(y))\,\|\nabla_Y v(y)\|)dy \\ \text{s.t.}\ & \int_X u(x)dx + \int_Y v(y)dy \geq (1-p_0)(\text{area}(X) + \text{area}(Y)). \end{aligned} \tag{18}$$

We can therefore solve (16) by alternatingly solving (17) and (18), which is expressed in the following framework algorithm:

1. *Weighted rigid alignment:* Fix u and v, and compute R and t minimizing (17).
2. Update $Y = RY + t$, and compute the misalignment fields $e_{X,Y}$ and $e_{Y,X}$.
3. *Optimal part selection:* Fix $e_{X,Y}$ and $e_{Y,X}$, and find u and v minimizing (18).

In what follows, we are going to present a discretization and a numerical scheme for this algorithm.

4 Numerical Framework

We assume the shape X to be given as a triangular mesh with N vertices $\{\mathbf{x}_1, ..., \mathbf{x}_N\} \subset \mathbb{R}^3$ represented as an $N \times 3$ matrix $\mathbf{X}$, and T faces, represented as a $T \times 3$ matrix $\mathbf{T}$ of vertex indices. For brevity, we will continue writing X referring to the mesh $(\mathbf{X}, \mathbf{T})$. We denote by $\mathbf{P}$ be the sparse $T \times N$ matrix, whose elements p_{ij} are set to $\frac{1}{3}$ if triangle i shares the vertex $\mathbf{x}_j$, and 0 otherwise. The matrix $\mathbf{P}$ can be thought of as a projection operator, converting a function defined on the mesh vertices into a function defined on the mesh faces. We denote by $\mathbf{a}'$ the $T \times 1$ vector, whose elements are areas of the mesh triangles. The $N \times 1$ vector and by $\mathbf{a} = \mathbf{P}\mathbf{a}'$ discretizes the area elements at each vertex on the mesh X. The membership function u is represented as the $N \times 1$ vector $\mathbf{u} = (u_1, ..., u_N)^{\mathrm{T}}$. The shape Y is represented similarly as a mesh $Y = (\mathbf{Y}, \mathbf{S})$, containing M vertices $\{\mathbf{y}_1, ..., \mathbf{y}_M\}$ and S faces. We denote by $\mathbf{Y}$, $\mathbf{S}$, $\mathbf{Q}$, $\mathbf{b}$, and $\mathbf{v}$ the counterparts of $\mathbf{X}$, $\mathbf{T}$, $\mathbf{P}$, $\mathbf{a}$, and $\mathbf{u}$, respectively.

4.1 Weighted Rigid Alignment

The solution of (17) in Step 1 of our framework algorithm is carried out using a symmetric weighted ICP algorithm. We fix $\mathbf{u}$ and $\mathbf{v}$, and iterate the following steps until convergence:

1. Compute the closest points

$$\mathbf{x}_i^* = \arg\min_{\mathbf{x}\in\{\mathbf{x}_1,...,\mathbf{x}_N\}} \|\mathbf{y}_i - \mathbf{x}\|; \qquad \mathbf{y}_i^* = \arg\min_{\mathbf{y}\in\{\mathbf{y}_1,...,\mathbf{y}_M\}} \|\mathbf{x}_i - \mathbf{y}\|, \tag{19}$$

   and construct the $3 \times M$ matrix $\mathbf{X}^*$ and the $3 \times N$ matrix $\mathbf{Y}^*$, whose rows are the $\mathbf{x}_i^*$, and the $\mathbf{y}_i^*$, respectively.

2. Compute the weighted centroids of the matrices $(\mathbf{X}, \mathbf{X}^*)$ and $(\mathbf{Y}^*, \mathbf{Y})$,

$$\overline{\mathbf{x}} = \sum_{i=1}^{N} u_i \mathbf{x}_i + \sum_{i=1}^{M} v_i \mathbf{x}_i^*, \qquad \overline{\mathbf{y}} = \sum_{i=1}^{N} u_i \mathbf{y}_i^* + \sum_{i=1}^{M} v_i \mathbf{y}_i, \tag{20}$$

   and construct the centered matrices $\overline{\mathbf{X}} = (\mathbf{x}_1 - \overline{\mathbf{x}}, ..., \mathbf{x}_M - \overline{\mathbf{x}}, \mathbf{x}_1^* - \overline{\mathbf{x}}, ..., \mathbf{x}_N^* - \overline{\mathbf{x}})$ and $\overline{\mathbf{Y}} = (\mathbf{y}_1^* - \overline{\mathbf{y}}, ..., \mathbf{y}_M^* - \overline{\mathbf{y}}, \mathbf{y}_1 - \overline{\mathbf{y}}, ..., \mathbf{y}_N - \overline{\mathbf{y}})$.
3. Compute the weighted covariance matrix $\mathbf{H} = \overline{\mathbf{X}}\mathbf{W}\overline{\mathbf{Y}}^{\mathrm{T}}$, where $\mathbf{W} = \mathrm{diag}(\mathbf{u}, \mathbf{v})$ is an $(N + M) \times (N + M)$ diagonal matrix with the elements of $\mathbf{u}$ and $\mathbf{v}$ on the diagonal. Compute the singular value decomposition $\mathbf{H} = \mathbf{U}\boldsymbol{\Lambda}\mathbf{V}^{\mathrm{T}}$, and set $\mathbf{R} = \mathbf{U}\mathbf{V}^{\mathrm{T}}$, and $\mathbf{t} = \overline{\mathbf{x}} - \mathbf{R}\overline{\mathbf{y}}$.
4. Replace the columns $\mathbf{y}_i$ of $\mathbf{Y}$ with their transformed versions $\mathbf{R}\mathbf{y}_i + \mathbf{t}$.

As a stopping condition, a norm of the change of $\mathbf{Y}$, e.g. $\sum \|\mathbf{y}_i - \mathbf{R}\mathbf{y}_i - \mathbf{t}\|^2 \leq \tau$ can be used for some small threshold τ. Once the ICP algorithm is terminated, we compute the local surface misalignment as

$$\begin{aligned} \mathbf{e}_{X,Y} &= (\|\mathbf{x}_1 - \mathbf{y}_1^*\|^2, ..., \|\mathbf{x}_N - \mathbf{y}_N^*\|^2)^{\mathrm{T}}, \\ \mathbf{e}_{Y,X} &= (\|\mathbf{y}_1 - \mathbf{x}_1^*\|^2, ..., \|\mathbf{y}_M - \mathbf{x}_M^*\|^2)^{\mathrm{T}}. \end{aligned} \tag{21}$$

Note that like in the classical ICP, when $e_{X,Y}(x) = \|x - y^*(x)\|^2$ and $e_{Y,X}(y) = \|y - x^*(y)\|^2$ are used, the best alignment between $\tilde{X}$ and $\tilde{Y}$ for a given correspondence is given analytically using the SVD of the covariance matrix $\mathbf{H}$.

4.2 Optimal Part Selection

The solution of (18) in Step 3 of our framework algorithm consists of finding the best parts $\mathbf{u}$ and $\mathbf{v}$ given fixed $\mathbf{e}_{X,Y}$ and $\mathbf{e}_{Y,X}$. We discretize the dissimilarity term $d(\tilde{X}, \tilde{Y})$ as $d(\mathbf{u}, \mathbf{v}) = \mathbf{e}_{X,Y}^{\mathrm{T}}\mathbf{A}\mathbf{u} + \mathbf{e}_{Y,X}^{\mathrm{T}}\mathbf{B}\mathbf{v}$, where $\mathbf{A} = \mathrm{diag}(\mathbf{a})$ and $\mathbf{B} = \mathrm{diag}(\mathbf{b})$ are, respectively, the $N \times N$ and $M \times M$ diagonal matrices with the elements of $\mathbf{a}$ and $\mathbf{b}$ on the diagonal. Arranging $\mathbf{e}_{X,Y}$ and $\mathbf{e}_{Y,X}$ into an $(M + N) \times 1$ vector $\mathbf{e}^{\mathrm{T}} = (\mathbf{e}_{X,Y}^{\mathrm{T}}, \mathbf{e}_{Y,X}^{\mathrm{T}})$, and $\mathbf{u}$ and $\mathbf{v}$ into the vector $\mathbf{w}^{\mathrm{T}} = (\mathbf{u}^{\mathrm{T}}, \mathbf{v}^{\mathrm{T}})$, we can write $d(\mathbf{w}) = \mathbf{e}^{\mathrm{T}}\mathbf{C}\mathbf{w}$, where $\mathbf{C} = \mathrm{diag}(\mathbf{a}, \mathbf{b})$ is the $(M + N) \times (M + N)$ diagonal matrix, containing the elements of $\mathbf{a}$ and $\mathbf{b}$ on the diagonal.

In order to discretize the irregularity term $r(\tilde{X})$, we first need to approximate the norm of the intrinsic gradient $\nabla_X u$ on the mesh X. Assuming a first-order approximation of u, the gradient $\nabla_X u$ is constant one each face of the mesh. Given a triangle i formed by the vertices $\mathbf{x}_{t_{i,1}}, \mathbf{x}_{t_{i,2}}, \mathbf{x}_{t_{i,3}}$, the gradient norm can be expressed as $g_i = \sqrt{\boldsymbol{\delta}^{\mathrm{T}}(\mathbf{X}_i^{\mathrm{T}}\mathbf{X}_i)^{-1}\boldsymbol{\delta}}$, where $\mathbf{X}_i = (\mathbf{x}_{t_{i,2}} - \mathbf{x}_{t_{i,1}}, \mathbf{x}_{t_{i,3}} - \mathbf{x}_{t_{i,1}})$ is a 3×2 matrix with the local system of coordinates of triangle t, and $\boldsymbol{\delta} = (u_{t_{i,2}} - u_{t_{i,1}}, u_{t_{i,3}} - u_{t_{i,1}})$ is the vector of the membership function differences. Arranging the g_i's as the elements of the $T \times 1$ vector $\mathbf{g}_X(\mathbf{u})$, we can express the irregularity term as $r(\mathbf{u}) = h(\mathbf{u})^{\mathrm{T}}\mathbf{A}\mathbf{P}\mathbf{g}_X(\mathbf{u})$, where $h(\mathbf{u})^{\mathrm{T}} = (h(u_1), ..., h(u_N))$ is the non-linearity h applied element-wise to the vector $\mathbf{u}$. The second irregularity term $r(\mathbf{v})$ is discretized in a similar way, yielding $r(\mathbf{v}) = h(\mathbf{v})^{\mathrm{T}}\mathbf{B}\mathbf{Q}\mathbf{g}_Y(\mathbf{v})$. Since the gradients $\mathbf{g}_X(\mathbf{u})$ and $\mathbf{g}_Y(\mathbf{v})$ are intrinsic, they are

invariant to rigid motion of X and Y and depend only on $\mathbf{u}$ and $\mathbf{v}$. In terms of the combined vector $\mathbf{w}$, the regularity $r(\tilde{X}) + r(\tilde{Y})$ can be written as

$$r(\mathbf{w}) = (h(\mathbf{u})^{\mathrm{T}}, h(\mathbf{v})^{\mathrm{T}}) \begin{pmatrix} \mathbf{AP} & \\ & \mathbf{BQ} \end{pmatrix} \begin{pmatrix} \mathbf{g}_X(\mathbf{u}) \\ \mathbf{g}_Y(\mathbf{v}) \end{pmatrix} = h(\mathbf{w})^{\mathrm{T}} \mathbf{D}\mathbf{g}(\mathbf{w}), \quad (22)$$

where $\mathbf{D} = \mathbf{C}\,\mathrm{diag}(\mathbf{P}, \mathbf{Q})$ and $\mathbf{g}(\mathbf{w})^{\mathrm{T}} = (\mathbf{g}_X(\mathbf{u})^{\mathrm{T}}, \mathbf{g}_Y(\mathbf{v})^{\mathrm{T}})$.

The partiality term $p(\tilde{X})$ is discretized as the inner product $p(\mathbf{u}) = \mathbf{a}^{\mathrm{T}}(\mathbf{1} - \mathbf{u})$, where $\mathbf{1}$ is an $N \times 1$ vector of ones. The second partiality term is discretized in a similar manner as $p(\mathbf{v}) = \mathbf{b}^{\mathrm{T}}(\mathbf{1} - \mathbf{v})$. Using these expressions, the inequality constraint in minimization problem (17) can be rewritten as the linear constraint $\mathbf{a}^{\mathrm{T}}\mathbf{u} + \mathbf{b}^{\mathrm{T}}\mathbf{v} \geq (1 - p_0)\alpha$, where $\alpha = \mathbf{a}^{\mathrm{T}}\mathbf{1} + \mathbf{b}^{\mathrm{T}}\mathbf{1}$ is the sum of the areas of the meshes X and Y. Using the combined vector $\mathbf{w}$, we can write $\mathbf{c}^{\mathrm{T}}\mathbf{w} \geq (1 - p_0)\alpha$ with $\mathbf{c}^{\mathrm{T}} = (\mathbf{a}^{\mathrm{T}}, \mathbf{b}^{\mathrm{T}})$.

Plugging the discretized dissimilarity, partiality and irregularity terms into problem (18), we obtain the

$$\min_{\mathbf{w}} \mathbf{e}^{\mathrm{T}}\mathbf{C}\mathbf{w} + \mu\, h(\mathbf{w})^{\mathrm{T}}\mathbf{D}\mathbf{g}(\mathbf{w}) \text{ s.t. } \begin{cases} \mathbf{c}^{\mathrm{T}}\mathbf{w} \geq (1 - p_0)\alpha \\ \mathbf{0} \leq \mathbf{w} \leq \mathbf{1} \end{cases} \quad (23)$$

Since the objective function has a well-structured sparse Hessian, the use of second-order minimization algorithms is appealing. In our implementation, the problem is solved using a reduced Newton descent algorithm [14], where the projection enforces the bound constraints $\mathbf{0} \leq \mathbf{w} \leq \mathbf{1}$. The inequality constraint $\mathbf{c}^{\mathrm{T}}\mathbf{w} \geq (1 - p_0)\alpha$ is introduced by adding an augmented Lagrangian term to the objective function. The reduced Newton system is solved using the modified Cholesky factorization, guaranteeing a descent direction at each iteration. Since part regularization has a low-pass filtering effect, the solution of (18) at low mesh resolution usually provides a sufficiently good initialization for the following resolution levels. Our practice shows that about four resolution levels improve convergence by about one order of magnitude.

The solution of the regularized matching problem yields the vector $\mathbf{w}$, representing the fuzzy part membership functions, and the transformation $\mathbf{R}$, $\mathbf{t}$ bringing those parts into best alignment. Since generally $\mathbf{w}$ will contain a continuous range of values from 0 to 1, the final step of the algorithm requires to "defuzzify" the fuzzy parts, i.e. convert $\mathbf{u}$ and $\mathbf{v}$ into binary membership functions. This is done by finding a threshold τ satisfying

$$\tau = \arg\min_{\tau} \tau \quad \text{s.t.} \quad \mathbf{c}^{\mathrm{T}}\boldsymbol{\chi}(\mathbf{w} \geq \tau) \geq (1 - p_0)\alpha, \quad (24)$$

where $\boldsymbol{\chi}(\mathbf{w} \geq \tau) = (\chi_1, ..., \chi_{M+N})^{\mathrm{T}}$ is the indicator vector with the elements $\chi_i = 1$ for $w_i \geq \tau$, and $\chi_i = 0$ otherwise. The crisp parts are then given by $\mathbf{u} = \boldsymbol{\chi}(\mathbf{u} \geq \tau)$ and $\mathbf{v} = \boldsymbol{\chi}(\mathbf{v} \geq \tau)$.

5 Results

In order to visualize the advantage of the proposed regularized matching framework, we compared it to unregularized matching performed by means of ICP with rejection of points based on distance and angle between corresponding normal vectors. Figures 2

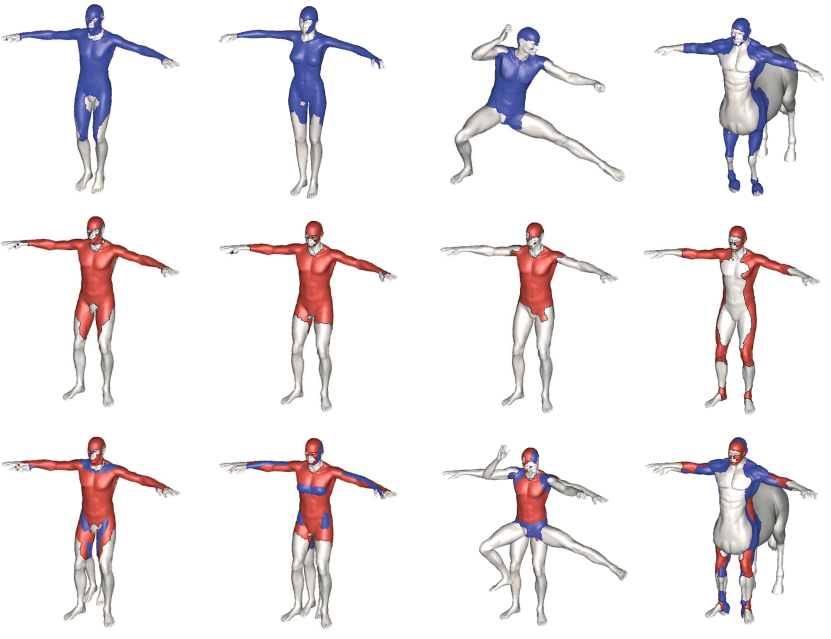

Fig. 2. Best partial matching of four rigid objects without regularization. Partiality p_0 is set to 0.4, 0.4, 0.65 and 0.65 (left-to-right). Two topmost rows present the best matching parts; bottom row shows their alignment. Highlighted are the matching parts.

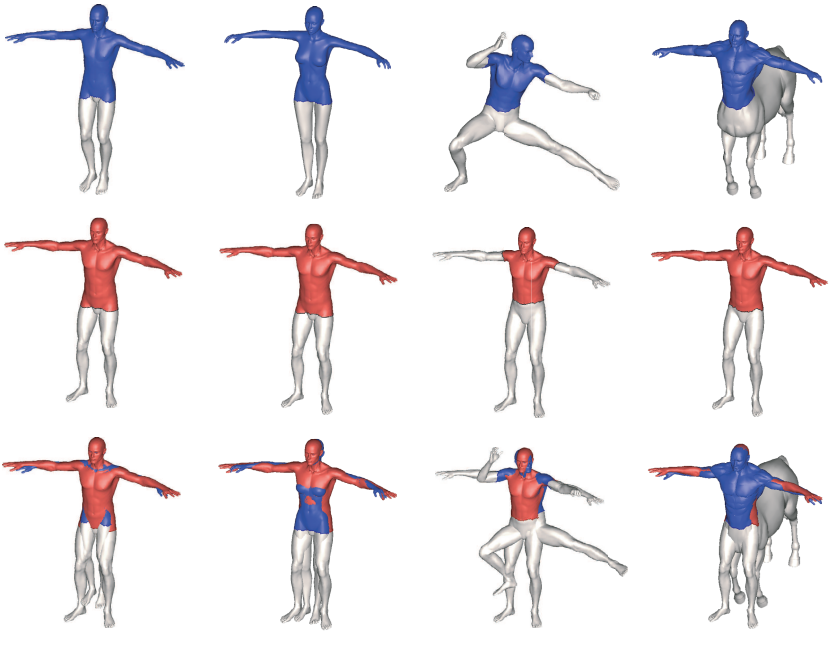

Fig. 3. Best partial matching of four rigid objects from Figure 2 with regularization

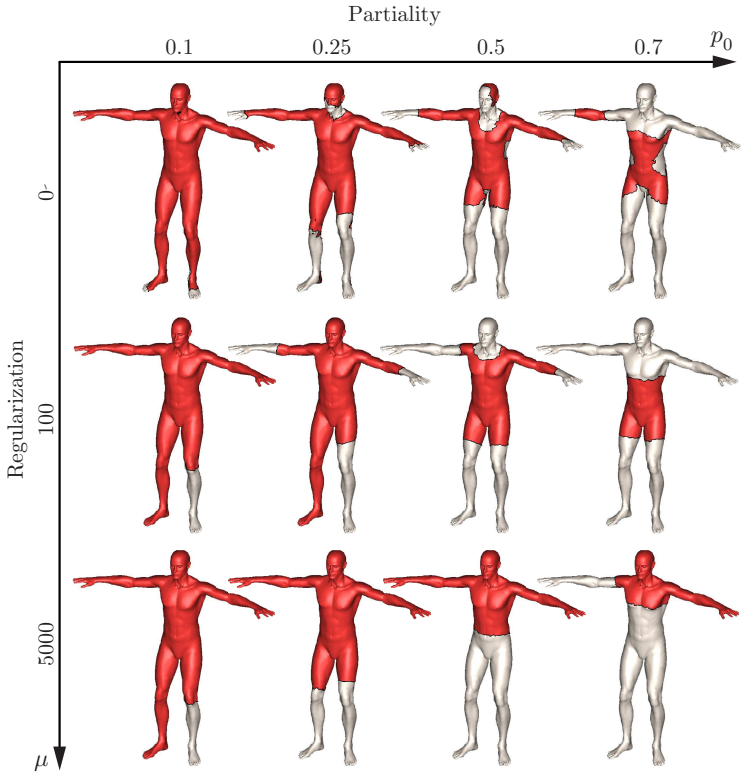

Fig. 4. Best partial matching of the male shape with the female shape from Figure 2 for different values of p_0 and μ. Note how for a fixed value of p_0, the area of the part remains the same, while its boundary length decreases with the increase of μ, resulting in a more regular part (see e.g. the third column).

and 3 show partial matching of four different rigid shapes with and without regularization, respectively. The use of regularization produces meaningful parts, unlike the unregularized case, where the parts contain many spurious disconnected components and holes. By changing the parameters p_0 and μ, we can control the size and the regularity of the selected parts.

Figure 4 shows the optimal part for different values of partiality p_0 and regularization parameter μ. Fixing p_0 and increasing μ gradually decreases the boundary length of the part, making it more regular.

6 Conclusions

In this paper, we considered the Paretian approach for partial matching of shapes proposed in [4] and based on a multicriterion problem of simultaneous maximization of similarity and size of the parts. We extended this approach, proposing a different

definition of part significance, taking into account the regularity of parts besides their size. We showed an efficient computation scheme based on fuzzy approximation, which allowed formulating the regularization in the spirit of the Mumford-Shah functional [13]. Though down-weighting of outliers is common in data fitting (including ICP algorithms), we believe that our paper proposes for the first time a rigorous method for selecting the weights optimally, in a controllable manner, and driven by the shape geometry, rather than resorting to heuristics. In our future studies, we intend to explore other definitions of topological and geometric regularity, and extensions of the proposed framework to nonrigid shapes.

References

1. Basri, R., Costa, L., Geiger, D., Jacobs, D.: Determining the similarity of deformable shapes. Vision Research 38, 2365–2385 (1998)
2. Besl, P.J., McKay, N.D.: A method for registration of 3D shapes. Trans. PAMI 14, 239–256 (1992)
3. Boiman, O., Irani, M.: Similarity by composition. In: Proc. NIPS (2006)
4. Bronstein, A.M., Bronstein, M.M., Bruckstein, A.M., Kimmel, R.: Paretian similarity of non-rigid objects. In: Sgallari, F., Murli, A., Paragios, N. (eds.) SSVM 2007. LNCS, vol. 4485, pp. 264–275. Springer, Heidelberg (2007)
5. Chan, T.F., Vese, L.A.: Active contours without edges. Trans. Image Processing 10(2), 266–277 (2001)
6. Chen, Y., Medioni, G.: Object modeling by registration of multiple range images. In: Proc. Conf. Robotics and Automation (1991)
7. Funkhouser, T., Shilane, P.: Partial matching of 3D shapes with priority-driven search. In: Proc. SGP (2006)
8. Gelfand, N., Mitra, N.J., Guibas, L., Pottmann, H.: Robust global registration. In: Proc. Symp. Geometry Processing (SGP) (2005)
9. Jacobs, D., Weinshall, D., Gdalyahu, Y.: Class representation and image retrieval with non-metric distances. Trans. PAMI 22, 583–600 (2000)
10. Kupeev, K., Wolfson, H.: On shape similarity. In: Proc. ICPR, pp. 227–237 (1994)
11. Latecki, L.J., Lakaemper, R., Wolter, D.: Optimal Partial Shape Similarity. Image and Vision Computing 23, 227–236 (2005)
12. Mitra, N.J., Gelfand, N., Pottmann, H., Guibas, L.: Registration of point cloud data from a geometric optimization perspective. In: Proc. Eurographics Symposium on Geometry Processing, pp. 23–32 (2004)
13. Mumford, D., Shah, J.: Boundary Detection by Minimizing Functionals. Image Understanding (1990)
14. Nocedal, J., Wright, S.: Numerical optimization. Springer, New York (1999)
15. Pentland, A.: Recognition by parts. In: Proc. ICCV, pp. 612–620 (1987)
16. Siddiqi, K., Kimia, B.: A shock grammar for recognition. In: Proc. CVPR, pp. 507–513 (1996)
17. Tangelder, J.W.H., Veltkamp, R.C.: A survey of content based 3D shape retrieval methods. In: Proc. Shape Modeling and Applications, pp. 145–156 (2004)