Not only size matters: regularized partial matching of nonrigid shapes

Alexander M. Bronstein
Michael M. Bronstein
Department of Computer Science, Technion – Israel Institute of Technology, Haifa 32000, Israel
{bron,mbron}@cs.technion.ac.il

Abstract

Partial matching is probably one of the most challenging problems in nonrigid shape analysis. The problem consists of matching similar parts of shapes that are dissimilar on the whole and can assume different forms by undergoing nonrigid deformations. Conceptually, two shapes can be considered partially matching if they have significant similar parts, with the simplest definition of significance being the size of the parts. Thus, partial matching can be defined as a multcriterion optimization problem trying to simultaneously maximize the similarity and the size of these parts. In this paper, we propose a different definition of significance, taking into account the regularity of parts besides their size. The regularity term proposed here is similar to the spirit of the Mumford-Shah functional. Numerical experiments show that the regularized partial matching produces semantically better results compared to the non-regularized one.

1. Introduction

Partial matching is probably one of the most challenging problems in nonrigid shape analysis. The problem consists of matching similar parts of shapes that are dissimilar on the whole and can assume different forms by undergoing nonrigid deformations. Such problems often arise in computer vision, for example, when the data to be matched are not available entirely due to acquisition imperfections. Depending on applications, partial matching can be used either to determine *partial similarity* between the shapes (a "distance" quantifying how different the shapes are) or *partial correspondence* (a relation between the points of the shapes). Typically, similarity problems are encountered in computer vision and pattern recognition, while correspondence problems are more often required in the realm of computer graphics.

Conceptually, two shapes can be considered partially matching if they have significant similar parts. For example, a centaur and a horse have a similar part (the horse body), which makes them partially similar [5]. Thus, partial matching can be found by segmenting the shape into significant parts and trying to match these parts separately [10, 7, 2]. However, "significance" is a semantic notion, and thus automatically finding such parts is not a well-defined problem. Many heuristic methods have been proposed in the literature for shape decomposition (see e.g. [1]). In shape retrieval applications, shapes are represented as collections of local descriptors, analogous to "bags of words" in text search engines. The significance of each local descriptor is determined statistically by its frequency [11]. Latecki *et al.* [8] performed partial matching by simplifying the shapes until they become the most similar. In [3], Bronstein *et al.* proposed solving a multicriterion optimization problem trying to simultaneously maximize the similarity and the size of the parts. It was argued that the size of the part is related to its significance: the larger is the part, the more significant it is.

In this paper, we argue that not only size matters. We propose a different definition of significance, taking into account the regularity of parts besides their size. The regularity term proposed here is similar to the spirit of the Mumford-Shah [9] functional, and can be considered as an extension thereof to non-Euclidean manifolds. The paper is organized as follows. In Section 2, we present the generic Paretian formulation of partial matching and the proposed regularization approach. In Section 3, we present a specific way of performing partial matching of nonrigid shapes and in Section 4 describe its numerical computation. Section 5 shows experimental results. Finally, Section 6 concludes the paper.

2. Regularized partial matching of shapes

Let X and Y be two shapes we would like to compare. We say that X and Y are *partially matching* if there exist *parts* $X' \subseteq X$ and $Y' \subseteq Y$ which are *similar* and *significant*. The degree of *dissimilarity* of parts can be expressed by a non-negative function $d : \Sigma_X \times \Sigma_Y \to \mathbb{R}_+$ (here Σ_X and Σ_Y denote the collection of all the parts of the shapes). As the measure of insignificance, Bronstein *et al.* [3] used the *partiality* function $p(X') = \mathrm{area}(X) - \mathrm{area}(X')$. In

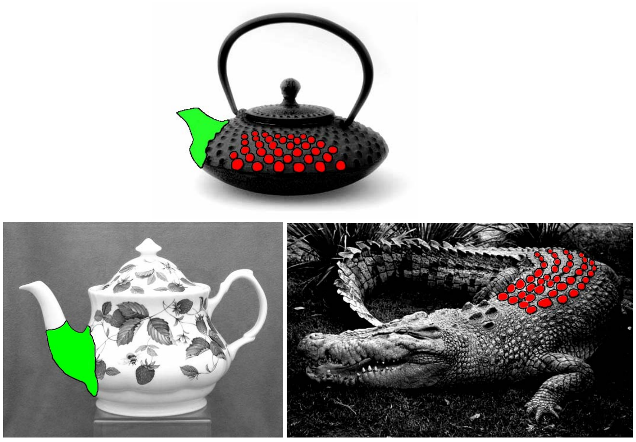

Figure 1. Conceptual illustration of an extreme case in which partial matching without regularization produces small disconnected parts and is semantically wrong (right). Partial matching taking into consideration not only the size of the parts but also their regularity produces semantically correct results (left).

this formulation, partial matching can be stated as the problem of simultaneous minimization of d and p over pairs of all the possible parts,

$$\min_{\Sigma_X \times \Sigma_Y} (d(X', Y'), p(X') + p(Y')). \quad (1)$$

A solution of the multicriterion optimization problem (1) is the set of parts (X^*, Y^*) achieving an optimal tradeoff between the dissimilarity and the partiality, in the sense that there exists no other pair of parts (X', Y') with both $d(X', Y') < d(X^*, Y^*)$ and $p(X') + p(Y') < p(X^*) + p(Y^*)$. Such a solution is called *Pareto optimal* and is not unique; it can be visualized as a curve in the (d, p) plane (referred to as the *Pareto frontier*) shown in bold in Figure 2. The value of d at the point with partiality p_0 on the Pareto frontier can be computed as

$$\min_{\Sigma_X \times \Sigma_Y} d(X', Y') \ \text{s.t.} \ p(X') + p(Y') \leq p_0. \quad (2)$$

In many cases, the solution of problem (2) manifests a tendency of finding multiple disconnected parts. An extreme case is visualized in Figure 1 (right): the bumps on a cast iron teapot and those on the back of an alligator make these two objects partially similar. While the area of the similar parts in this example is large, the parts are small and fragmented. Semantically, one large and regular part (even of smaller area) shown in Figure 1 (left) is preferable. This conceptual example is a motivation to the main idea of this paper: when the significance of parts in the partial similarity problem is considered, not only the size matters. Quantitatively, we can measure the "quality" of the part X' using some *irregularity* function $r(X')$. We are thus looking for the largest, most similar and most regular parts, giving rise to the following multicriterion optimization problem,

$$\min_{\Sigma_X \times \Sigma_Y} (d(X', Y'), p(X') + p(Y'), r(X') + r(Y')), \quad (3)$$

The set of the Pareto optimal solutions of problem (3) can be visualized as a three-dimensional surface in the (d, p, r) space (Figure 2). The new formulation can also be regarded as a regularized version of problem (2),

$$\min_{\Sigma_X \times \Sigma_Y} d(X', Y') + \mu(r(X') + r(Y')) \quad \text{s.t.} \ p(X') + p(Y') \leq p_0, \quad (4)$$

which we call here the *regularized partial matching*. For non-zero μ, the regularity term will give preference to parts with larger regularity even at the cost of smaller area or larger dissimilarity. Alternatively, we can rewrite (4) as an unconstrained minimization problem, and interpret the aggregate $p(X') + p(Y') + \mu(r(X') + r(Y'))$ as a new definition of significance, taking into consideration not only the size but also the regularity of the parts.

3. Partial matching of nonrigid shapes

In this section, we show specific definitions used in the regularized partial matching problem of nonrigid shapes. Following [3], we model a nonrigid shape as a compact *metric space* (X, d_X), where d_X is the *geodesic metric*, measuring the lengths of the shortest paths between pairs of points on the shape. Geometric quantities expressed in terms of d_X are referred to as *intrinsic*, as they do not depend on the way the shape is laid out in the ambient Euclidean space. Such quantities are invariant under inelastic deformations of the shape. Two nonrigid shapes (X, d_X) and (Y, d_Y) are said to be *isometric* if there exists a bijective map $\varphi : X \to Y$ such that $d_X(x, x') = d_Y(\varphi(x), \varphi(x'))$ for all x, x' in X. Since in practice the deformations of nonrigid shapes are rarely truly isometric, we can relax the requirement of φ being metric-preserving and require that it does not distort the metric significantly. The *distortion* of the metric can be quantified, for example, as

$$\begin{aligned} \text{dis}(\varphi) &= \int_{X \times X} |d_X(x, x') - d_Y(\varphi(x), \varphi(x'))|^2 dx dx' \\ &= \int_{X \times X} e_\varphi(x, x') dx dx', \end{aligned} \quad (5)$$

where $e_\varphi(x, x')$ measures the "local" distortion of φ at each pair of points (x, x') on X. Using the distortion as a criterion of dissimilarity, we define the *intrinsic dissimilarity* of the shapes X and Y as

$$d(X, Y) = \min_{\varphi : X \to Y} \text{dis}(\varphi) + \min_{\psi : Y \to X} \text{dis}(\psi). \quad (6)$$

Since our final goal is partial shape matching, we need to extend the above definition to comparison of parts. For this purpose, we define a *part* of the shape X as a metric sub-space $(X' \subseteq X, d_X|_{X'})$, where $d_X|_{X'}$ is the *restricted metric*, equal to d_X on $X' \times X'$. The distance between parts is measured as

$$\begin{aligned} d(X',Y') &= \min_{\varphi:X'\to Y'} \mathrm{dis}(\varphi) + \min_{\psi:Y'\to X'} \mathrm{dis}(\psi) \quad (7)\\ &= \min_{\varphi:X'\to Y'} \int_{X'\times X'} e_\varphi(x,x')dxdx' \\ &+ \min_{\psi:Y'\to X'} \int_{Y'\times Y'} e_\psi(y,y')dydy'. \end{aligned}$$

In addition to the function d measuring the dissimilarity of parts, minimization problems (3) and (4) require the partiality p and irregularity r to be defined. We define the *partiality* of a part as the area of $X'^c = X \setminus X'$,

$$p(X') = \int_X dx - \int_{X'} dx. \quad (8)$$

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

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

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*. Other irregularity criteria can be used as well to suit the specific application needs. In the sequel, we will use the above irregularity for the sake of simplicity.

3.1. Fuzzy formulation

The solution of (3) or (4) 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 *fuzzy* approximation $u : X \to [0,1]$ and $v : Y \to [0,1]$, respectively, bringing the problem back to a tractable continuous formulation.

The dissimilarity of fuzzy parts u and v can be expressed analogously to (7) by weighting the integrand with the respective fuzzy membership functions,

$$\begin{aligned} d(u,v) = \min_{\varphi:X\to Y} &\int_{X\times X} u(x)u(x')e_\varphi(x,x')dxdx' \\ + \min_{\psi:Y\to X} &\int_{Y\times Y} v(y)v(y')e_\psi(y,y')dydy'. \quad (10) \end{aligned}$$

Since $d(u,v)$ can be decoupled into two independent minimization problems with respect to φ and ψ, we will denote by $d(u)$ and $d(v)$ the first and the second terms of $d(u,v)$, respectively, and write $d(u,v) = d(u) + d(v)$.

The fuzzy version of the partiality can be written as

$$p(u) = \int_X (1-u(x))dx. \quad (11)$$

The irregularity term is slightly more elaborate, since in the fuzzy formulation there is no "boundary". However, adopting the spirit of the Mumford-Shah approach, we can replace integration along the boundary by integration of the band in which the membership function changes from small to large values [9],

$$r(u) = \int_X h(u(x)) \|\nabla_X u(x)\| \, dx, \quad (12)$$

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.

Plugging the former expressions into the partial shape matching problem (4), we obtain

$$\begin{aligned} &\min_{u,v,\varphi,\psi} \quad (13)\\ &\int_X \left(\int_X u(x)u(x')e_\varphi(x,x')dx' + \mu h(u(x)) \|\nabla_X u(x)\| \right) dx \\ &\int_Y \left(\int_Y v(y)v(y')e_\psi(y,y')dy' + \mu h(v(y)) \|\nabla_Y v(y)\| \right) dy \\ &\text{s.t.} \int_X u(x)dx + \int_Y v(y)dy \geq \mathrm{area}(X) + \mathrm{area}(Y) - p_0. \end{aligned}$$

Splitting the optimization variables into u, v and φ, ψ, problem (14) can be solved using the following alternating minimization algorithm:

1. Fix u and v and find φ and ψ by solving

$$\begin{aligned} \varphi &= \operatorname*{argmin}_{\varphi:X\to Y} \int_{X\times X} u(x)u(x')e_\varphi(x,x')dxdx'; \\ \psi &= \operatorname*{argmin}_{\psi:Y\to X} \int_{Y\times Y} v(y)v(y')e_\psi(y,y')dydy'. \end{aligned}$$

2. Fix e_φ and e_ψ and find u and v by solving

$$\begin{aligned} &\min_{u,v} \int_{X\times X} u(x)u(x')e_\varphi dxdx' + \int_{Y\times Y} v(y)v(y')e_\psi dydy' \\ &\int_X h(u(x)) \|\nabla_X u(x)\| \, dx + \int_Y h(v(y)) \|\nabla_Y v(y)\| \, dy \\ &\text{s.t.} \int_X u(x)dx + \int_Y v(y)dy \geq \mathrm{area}(X) + \mathrm{area}(Y) - p_0. \end{aligned}$$

3. Iterate Steps 1 – 2 until convergence.

The output of the algorithm are the optimal fuzzy parts u and v, the minimum-distortion correspondences φ and ψ, and the values of $d(u,v)$, $p(u)+p(v)$, and $r(u)+r(v)$. Thus, both partial similarity and correspondence are obtained by solving the same problem. We call the value of $d(u,v)$ at the solution *scalar partial similarity*. In the following section, we are going to present a discretization of this algorithm.

4. Discretization

We assume the input shapes to be given as triangular meshes, which for the sake of convenience we will continue denoting as X and Y. The mesh X comprises M vertices $\{x_1, ..., x_M\}$, whose coordinates in $\mathbb{R}^3$ are represented as an $M \times 3$ matrix $\mathbf{X}$, and T faces, represented as a $T \times 3$ matrix of vertex indices. Similarly, the mesh Y is represented by an $N \times 1$ matrix $\mathbf{Y}$ of vertices, and an $S \times 1$ matrix $\mathbf{S}$ of face indices. The fuzzy membership functions $u(x)$ and $v(y)$ are discretized on X and Y, and denoted by the vectors $\mathbf{u} = (u_1, ..., u_M)^{\mathrm{T}}$ and $\mathbf{v} = (v_1, ..., v_N)^{\mathrm{T}}$, respectively. The area elements on X are discretized and represented as the vector $\mathbf{a} = (a_1, ..., a_N)^{\mathrm{T}}$, where a_i is set to be $\frac{1}{3}$ of the sum of the areas of the triangular faces of X sharing the vertex x_i. The discrete area elements of Y are computed in the same way, and are denoted by $\mathbf{b} = (b_1, ..., b_N)^{\mathrm{T}}$.

We first discretize Step 1 in the alternating minimization algorithm from the previous section. Since the dissimilarity $d(u,v)$ in (10) can be decoupled into two independent terms $d(u)$ and $d(v)$, in what follows we will show the discretization of $d(u)$ on the mesh X only. The other term $d(v)$ is discretized on Y in the same manner. The mesh X is sampled at m points using the *farthest point sampling* procedure [4]. For simplicity, we assume this set of points to be a sub-set of the vertices of X, numbered without loss of generality as $X_m = \{x_1, ..., x_m\}$. We also assume to be given two linear operators, an $m \times M$ matrix $\mathbf{P}$, and an $M \times m$ matrix $\mathbf{P}'$, projecting a function on X onto X_m, and vice versa. To simplify our discussion, we will assume that $\mathbf{P}' = \mathbf{P}^{\mathrm{T}}$, though other possibilities to construct $\mathbf{P}$ and $\mathbf{P}'$ also exist. Projecting $\mathbf{u}$ onto X_m, we can write

$$d(\mathbf{u}) = \min_{\varphi: X_m \to Y} \sum_{i,j=1}^{m} u'_i u'_j |d_X(x_i,x_j) - d_Y(\varphi(x_i),\varphi(x_j))|^2 a'_i a'_j \quad (14)$$

where $\mathbf{u}' = \mathbf{P}\mathbf{u}$, and $\mathbf{a}' = \mathbf{P}\mathbf{a}$. Minimization over all mappings $\varphi : X_m \to Y$ can be reformulated in terms of the images $y'_i = \varphi(x_i)$ as

$$d(\mathbf{u}) = \min_{\{y'_1,...,y'_m\}} \sum_{i,j=1}^{m} u'_i u'_j |d_X(x_i,x_j) - d_Y(y'_i,y'_j)|^2 a'_i a'_j. \quad (15)$$

Minimization problem (15) is solved using the weighted *generalized multidimensional scaling* algorithm detailed in [3]. Since the x_i are fixed, distance terms $d_X(x_i,x_j)$ can be precomputed using, for example, the fast marching algorithm [6]. On the other hand, the terms $d_Y(y'_i,y'_j)$ involve optimization variables and are approximated using the geodesic interpolation procedure from [3].

Once $d(\mathbf{u})$ is computed, we construct the $m \times m$ matrix $\mathbf{E}'_\varphi$, whose elements are given by

$$e_i = |d_X(x_i,x_j) - d_Y(y'_i,y'_j)|^2. \quad (16)$$

Since $\mathbf{E}'_\varphi$ is a function on $X_m \times X_m$, we can use $\mathbf{P}^{\mathrm{T}}$ to project it onto $X \times X$, obtaining the $M \times M$ matrix $\mathbf{E}_\varphi = \mathbf{P}^{\mathrm{T}}\mathbf{E}'_\varphi\mathbf{P}$. The second dissimilarity term $d(\mathbf{v})$ and the $N \times N$ matrix $\mathbf{E}_\psi$ are computed in an analogous way.

Discretization of Step 2 in the alternating minimization algorithm requires the discretization of the dissimilarity terms $d(u)+d(v)$, the regularity terms $r(u)+r(v)$, and the constraint on partiality in minimization problem (14). The dissimilarity terms can be straightforwardly discretized as the quadratic forms

$$d(\mathbf{u}) = \mathbf{u}^{\mathrm{T}}\mathbf{A}\mathbf{E}_\varphi\mathbf{A}\mathbf{u}, \qquad d(\mathbf{v}) = \mathbf{v}^{\mathrm{T}}\mathbf{B}\mathbf{E}_\psi\mathbf{B}\mathbf{v}, \quad (17)$$

where $\mathbf{A} = \mathrm{diag}(\mathbf{a})$ and $\mathbf{B} = \mathrm{diag}(\mathbf{b})$ denote the diagonal matrices with the discrete area elements $\mathbf{a}$ and $\mathbf{b}$ on the diagonal, respectively.

In order to discretize the irregularity term $r(u)$, 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 $x_{t_{i,1}}, x_{t_{i,2}}, x_{t_{i,3}}$, 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 $\mathbf{X}_i = (x_{t_{i,2}} - x_{t_{i,1}}, x_{t_{i,3}} - x_{t_{i,1}})$ 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. We arrange the g_i's as the elements of the $T \times 1$ vector $\mathbf{g}_X(\mathbf{u})$. Since $\mathbf{g}_X$ is a function on the faces of X, while the membership $\mathbf{u}$ and the discrete area elements $\mathbf{a}$ are defined on the vertices, we need to project $\mathbf{g}_X$ onto X. For this purpose, we construct an $M \times T$ matrix $\mathbf{Q}$, whose elements q_{ij} are set to $\frac{1}{3}$ if triangle j shares the vertex x_i, and 0 otherwise. Using this projection operator, we can we can express the irregularity term as $r(\mathbf{u}) = h(\mathbf{u})^{\mathrm{T}}\mathbf{A}\mathbf{Q}\mathbf{g}_X(\mathbf{u})$, where $h(\mathbf{u})^{\mathrm{T}} = (h(u_1), ..., h(u_M))$ is the approximation of the delta function $h(t) \approx \delta(t-0.5)$ applied element-wise to the vector $\mathbf{u}$. The second irregularity term $r(\mathbf{v})$ is discretized in a similar way, yielding $r(\mathbf{v}) = h(\mathbf{v})^{\mathrm{T}}\mathbf{B}\mathbf{R}\mathbf{g}_Y(\mathbf{v})$, where $\mathbf{R}$ is the $N \times S$ analog of the projection matrix $\mathbf{Q}$. The partiality constraint in Step 2 can be straightforwardly discretized as $\mathbf{a}^{\mathrm{T}}\mathbf{u} + \mathbf{b}^{\mathrm{T}}\mathbf{v} \geq \mathbf{1}^{\mathrm{T}}\mathbf{a} + \mathbf{1}^{\mathrm{T}}\mathbf{b} - p_0$, where $\mathbf{1}$ denotes a vector of ones.

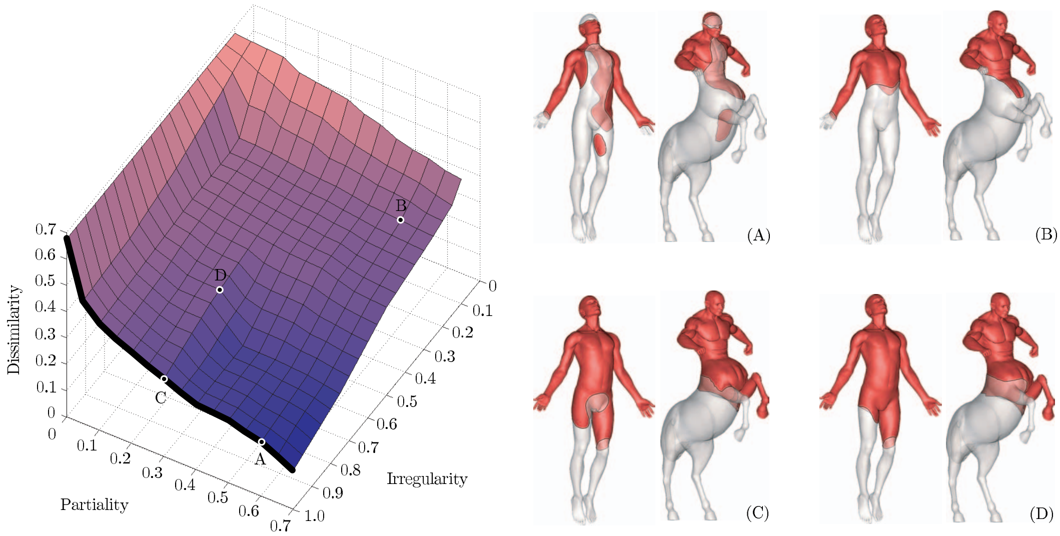

Figure 2. 3D Pareto frontier in the regularized partial matching problem. Shown in bold curve is the 2D Pareto frontier of the non-regularized problem. A, C: non-regularized parts; B, D: regularized parts.

Arranging $\mathbf{u}$ and $\mathbf{v}$ into an $(M+N)\times 1$ vector $\mathbf{w} = (\mathbf{u};\mathbf{v})$, and plugging in the former expressions, we can formulate a discrete version of Step 2 in the alternating minimization algorithm as the solution of

$$\min_{0\leq \mathbf{w}\leq 1} \mathbf{w}^{\mathrm{T}} \begin{pmatrix} \mathbf{AE}_{\varphi}\mathbf{A} & \\ & \mathbf{BE}_{\psi}\mathbf{B} \end{pmatrix} \mathbf{w} + \mu h(\mathbf{w})^{\mathrm{T}} \begin{pmatrix} \mathbf{AQ} & \\ & \mathbf{BR} \end{pmatrix} \mathbf{g}(\mathbf{w})$$
$$\text{s.t. } (\mathbf{a};\mathbf{b})^{\mathrm{T}}\mathbf{w} \geq (\mathbf{a};\mathbf{b})^{\mathrm{T}}\mathbf{1} - p_0 \quad (18)$$

where $\mathbf{g}(\mathbf{w}) = (\mathbf{g}_X(\mathbf{u}); \mathbf{g}_Y(\mathbf{v}))$.

5. Results

In this section, we show three experiments to exemplify the proposed method. The experiments were performed on objects from the Nonrigid world dataset available online at `http://tosca.technion.ac.il`.Each shape in the dataset was represented as a triangular mesh with 2000 vertices. Geodesic distances were measured using fast marching [6]. We used a BFGS quasi-Newton minimization algorithm for the solution of (18) in Step 2 of the alternating minimization algorithm. The minimum distortion correspondence (15) in Step 1 was computed using GMDS [3].

In the first experiment, we computed the 3D Pareto frontier representing the tradeoff between the dissimilarity, partiality and irregularity in the problem of partial matching of a centaur and a man (Figure 2). The results produced by the method of [3] can be considered as a particular case where no regularization is used, visualized by a bold curve in the figure. It can be clearly seen that the parts are fragmented (Figure 2 A and C). Regularization makes the resulting parts better and the matching more meaningful (Figure 2 B and D).

In the second experiment, we compared full similarity and partial similarity criteria on the Nonrigid world dataset. The dataset consisted of five objects (centaur, horse, seahorse, male and female); each object appeared in five different instances produced by near-isometric deformations. Full similarity of shapes was computed using GMDS. Partial similarity was computed using the presented approach. Figure 3 shows the confusion matrices representing the results. The full similarity criterion is insensitive to the intra-class variability of the shapes (i.e., different deformations of the same objects are similar). However, it fails to capture correctly the inter-class similarity: the centaur, horse and seahorse appear dissimilar. On the other hand, the partial similarity criterion captures correctly the similarity of the centaur, horse and the seahorse (Figure 3, bottom).

In the third experiment, regularized intrinsic partial correspondence between the shapes of a centaur, a horse and a man was computed. Figure 4 visualizes the correspondence by plotting the Voronoi regions around each of the corresponding points. The obtained correspondence is meaningful and accurate despite significant deformations of the shapes and large non-overlapping parts.

6. Conclusions

In this paper, we considered the Paretian approach for partial matching of shapes based on a multicriterion problem of simultaneous maximization of similarity and significance of the parts. In [3], the significance was measured

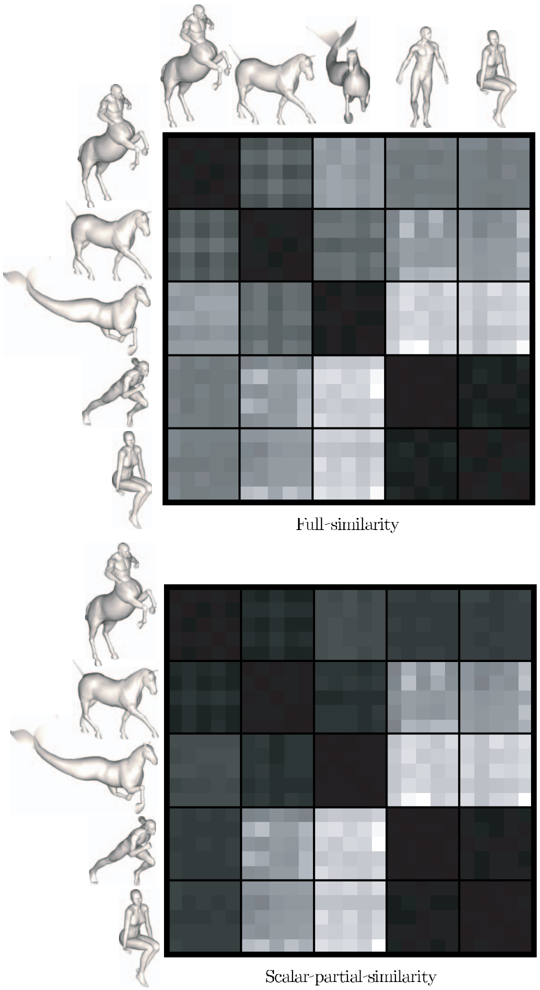

Figure 3. Confusion matrices representing full dissimilarity (top) and partial dissimilarity (bottom) between a set of nonrigid shapes. Darker shade of gray stands for higher similarity.

as the size of the parts. Here, we extended this approach, proposing a different definition of part significance, which takes 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 [9] functional. In our future studies, we intend to explore other definitions of topological and geometric regularity.

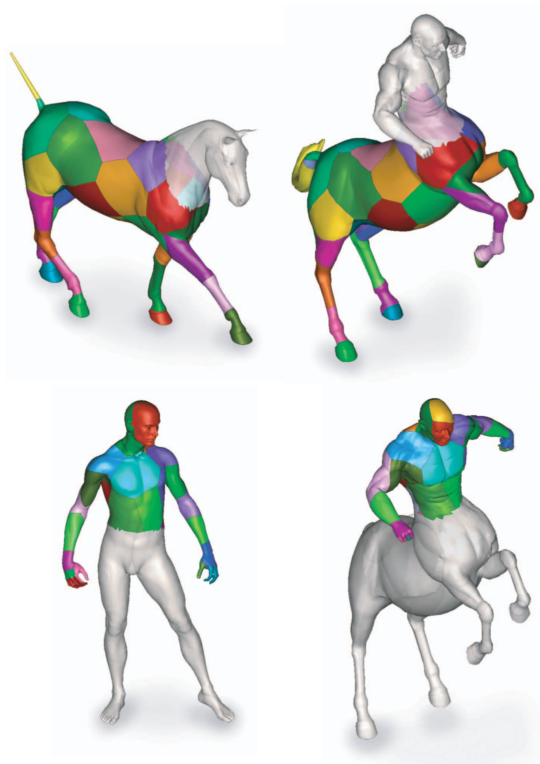

Figure 4. Partial correspondence between a horse and a centaur (left) and a centaur and a human (right), computed using the proposed algorithm.

References

[1] R. Basri, L. Costa, D. Geiger, and D. Jacobs. Determining the similarity of deformable shapes. *Vision Research*, 38:2365–2385, 1998. 1

[2] O. Boiman and M. Irani. Similarity by composition. In *Proc. NIPS*, 2006. 1

[3] A. M. Bronstein, M. M. Bronstein, A. M. Bruckstein, and R. Kimmel. Analysis of two-dimensional non-rigid shapes. *IJCV*. 1, 2, 4, 5

[4] T. F. Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoret. Comp. Sci.*, 38(2):293–306, 1985. 4

[5] D. Jacobs, D. Weinshall, and Y. Gdalyahu. Class representation and image retrieval with non-metric distances. *Trans. PAMI*, 22:583–600, 2000. 1

[6] R. Kimmel and J. A. Sethian. Computing geodesic paths on manifolds. *Proc. National Academy of Sciences (PNAS)*, 95(15):8431–8435, 1998. 4, 5

[7] K. Kupeev and H. Wolfson. On shape similarity. In *Proc. ICPR*, pages 227–237, 1994. 1

[8] L. J. Latecki, R. Lakaemper, and D. Wolter. Optimal Partial Shape Similarity. *Image and Vision Computing*, 23:227–236, 2005. 1

[9] D. Mumford and J. Shah. Boundary Detection by Minimizing Functionals. *Image Understanding*, 1990. 1, 3, 6

[10] A. Pentland. Recognition by parts. In *Proc. ICCV*, pages 612–620, 1987. 1

[11] J. W. H. Tangelder and R. C. Veltkamp. A survey of content based 3D shape retrieval methods. In *Proc. Shape Modeling Applications*, pages 145–156, 2004. 1