Measuring 3D Shape Similarity by Matching the Medial Scaffolds

Ming-Ching Chang
LEMS, Brown University
Providence RI 02912 USA
mcchang@lems.brown.edu

Benjamin B. Kimia
LEMS, Brown University
Providence RI 02912 USA
kimia@lems.brown.edu

Abstract

We propose to measure 3D shape similarity by matching a medial axis ($\mathcal{MA}$) based representation—the medial scaffold *($\mathcal{MS}$). Shape similarity is measured as the minimum extent of deformation necessary for one shape to match another, guided by the $\mathcal{MS}$. This approach is an extension of an approach to match 2D shapes by matching their* shock graphs*, whereas here in 3D the $\mathcal{MS}$ is in the form of a* hypergraph*. The $\mathcal{MS}$ representation is both hierarchical and complete. Our approach finds the* optimal *deformation path between two shapes by modelling shape deformations as discrete topological changes (*transitions*) of the $\mathcal{MS}$, with costs associated with each transition. We first regularize the $\mathcal{MS}$ hypergraphs and use the* graduated assignment *graph matching scheme to match the hypergraphs. A set of compatibility functions is defined to measure the pairwise similarity between the $\mathcal{MS}$ nodes, curves, and sheets. Early results on matching carpal bones and other shapes promise its potential in a range of applications.*

1. Introduction

Measuring 3D shape similarity is an important task in object recognition; applications include shape retrieval and clustering in databases [8], querying industrial parts, matching bio-chemical structures, *etc*. Central to this task is the issue of shape **representation**. Typically a *descriptor* is extracted from the shape, usually with a great deal of simplification to enable efficient matching. The choice of descriptor is often domain-specific and could vary largely from one application to another, facing the dilemma of either being too coarse (ignoring information) or too complex (redundant and unstable). As the problem of matching *rigid closed* shapes is generalized into matching *partial* or *articulated* shapes, developing a *generic* representation becomes significant and is the key of this paper.

A major branch in shape representation is the symmetry-based **medial axis** ($\mathcal{MA}$). The $\mathcal{MA}$ is promising for shape recognition [19, 21] in that *(i)* it organizes the shape information in a *hierarchical*, *intrinsic* graph-like structure [16], which enables matching parts of deformed shapes naturally,

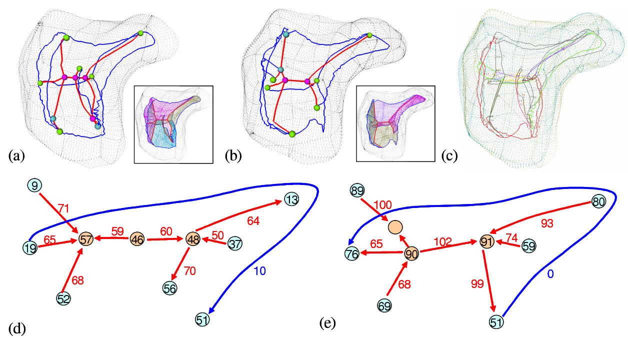

Figure 1. The matching of the $\mathcal{MS}$ hypergraphs of two carpal bones [17] in (a) and (b) is shown in (c). (d,e) show a manual correspondence, where the graph components are labeled in numbers, serving as the ground truth to validate the automatic matching.

and *(ii)* such information captured with the $\mathcal{MA}$ is *complete* in that a full shape reconstruction is always possible [10]. Despite these advantages, the $\mathcal{MA}$ is generally sensitive to perturbation and difficult to model in the 3D case. Such issues have been recently addressed [12, 11]. We adopt the **Medial Scaffold** ($\mathcal{MS}$)—a hierarchical organization of the 3D $\mathcal{MA}$ into a *hypergraph* form [16] and a *regularization* framework of the $\mathcal{MS}$ [4] to deal with the above barriers. The $\mathcal{MA}$ instabilities which induce sudden topological changes are formally classified as a set of *transitions* and thus can be regularized via a set of *transforms*. We propose to match the regularized $\mathcal{MS}$ such as the ones shown in Fig.1 to estimate a global similarity between shapes.

Our main contribution is a novel solution to measure 3D shape similarity by matching the $\mathcal{MS}$ hypergraphs representing the underlying shapes. Following a theoretical framework to measure shape similarity as the *minimum deformation* necessary for one shape to match the other in 2D [19], the matching here is guided by the $\mathcal{MS}$ as a representation, which retains both key benefits of the $\mathcal{MA}$ (*hierarchical* and *complete*). The amount of shape deformation can be formulated as an integration of *infinitesimal elastic* changes to optimally match the $\mathcal{MS}$ branches (sheets and curves). Fig.2a [3] illustrates an example. We propose to approximate this optimal solution by first regularizing the $\mathcal{MS}$ [4] and matching the $\mathcal{MS}$ hypergraphs [5, 13]. Two

2009 IEEE 12th International Conference on Computer Vision Workshops, ICCV Workshops
978-1-4244-4441-0/09/$25.00 ©2009 IEEE

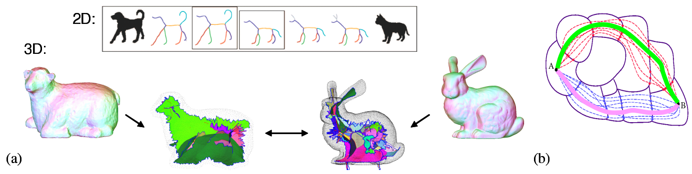

Figure 2. (a) The shock graph in 2D and the $\mathcal{MS}$ in 3D as the representation for finding the optimal deformation in matching two shapes. (b) Deformation path represented as a sequence of discrete shock *transitions*. Each blob represents a *shape cell* [19] where all shapes share a common shock/$\mathcal{MS}$ topology.

improvements are significant: *(i)* a natural extension of the graph matching scheme to match *hypergraphs* and *(ii)* a set of similarity measures to reflect the *structural* and *parametric* differences of the $\mathcal{MS}$ hypergraphs.

The paper is organized as follows. §2 reviews the background in 3D shape similarity matching. §3 describes the $\mathcal{MS}$ as our representation for 3D shapes. §4 covers our extended graph matching scheme, and §5 elaborates the compatibility measures between the $\mathcal{MS}$ hypergraph components (nodes, curves, sheets). After the matching is performed, §6 computes the final similarity measure by summing up the compatibility measures weighted by the assignment coefficients. Finally, our approach is examined in §7 on matching medical (carpal bones) and synthetic shapes.

2. Background

Measuring shape similarity in recognition is a fundamental problem with an abundant literature; refer to [2] for surveys in the 3D domain. We briefly organize recent approaches into two main categories, namely, the *(i)* shape *descriptor*-based and *(ii)* structural *graph*-based methods.

Descriptor-based methods are the current mainstream. A shape descriptor (feature, signature) is extracted to describe the shape and distinguish it from others. A large variety of descriptors have been proposed, which are briefly classified into *four* sub-categories: *(i)* local feature based, which relies on local salient geometric features [9, 15] such as the curvature or primitives of flat regions [14]; *(ii) spherical* functions, such as the spherical harmonic, shape histogram, and shape context; *(iii) statistical measure*-based, such as the shape distribution and other generalizations; *(iv) view*-based, by matching 2D views of the 3D objects; and *(v) voxel*-based, assuming the input is a solid volume where a distance transform can be effectively computed. We note that only a few methods are capable of handle *partial*, *non-closed*, and *sampled* shapes (as we do in this paper) such as the *spin-image* [15] and others [14].

Graph-based methods employ a graph to represent the connectivity between *parts* of a shape, where the matching of partial or deformed shapes can be carried out naturally. Recent works are organized by the type of graph used in the methods as follows: *(i)* the *Reeb graph* is a topological graph based on the Morse analysis on a pre-defined function (such as the geodesic or height functions), where graph edges are not necessarily skeletal/symmetric [1]; *(ii)* the *skeletal* graph such as [24] encodes topological signature vectors for matching; *(iii)* the *curve-skeleton* is an one-dimensional centerline roughly central inside the 3D shape. Although it is much simpler than the $\mathcal{MA}$ (which consists of 2D sheets), a suitable mathematical definition for such "centredness" still needs investigation [7]; *(iv)* two-dimensional medial *sheets*: Pizer *et al.* pioneer in using fixed-topology medial models for segmentation [23]; Siddiqi *et al.* [21] employ a directed acyclic graph of the *medial sheets* to retrieve articulated 3D models. Our approach belongs to this category. One significance of our method is an explicit exploit of the medial sheet *connectivity* [4] to extract a 3D structure for matching — a direct extension from a simplified 2D analysis in [19, 20].

3. Medial Scaffold ($\mathcal{MS}$) as Representation

The 3D medial axis ($\mathcal{MA}$) generally consists of medial sheets (A_1^2), curves (A_3, A_1^3), and nodes (A_1A_3, A_1^4) [11], where the A_k^n notation [11] indicates the order (k) of contact of a maximal ball with n surface points used to classify the $\mathcal{MA}$ points: An A_1^3 curve (red in Fig.1) delimits A_1^2 sheets at an *axis* and ends at A_1^4 or A_1A_3 nodes. An A_3 curve (blue) delimits an A_1^2 sheet at a *rib* and ends at A_1A_3 nodes. We represent the 3D $\mathcal{MA}$ as the *Medial Scaffold* ($\mathcal{MS}$) [16, 4], which as well organize the abundant information of shape into a hierarchical *hypergraph* form. The $\mathcal{MS}$ has several advantages in modeling 3D shapes: *(i)* Shape information is organized *intrinsically* with the $\mathcal{MS}$ structure and is *complete* (in allowing a full reconstruction of the shape). *(ii)* Instabilities of the $\mathcal{MA}$ can be formally handled as *transitions* [12], which are sudden $\mathcal{MA}$ topological changes corresponding to perturbation of shape. In 3D there are **7** generic *transitions*, which can be regularized case-by-case via **11** *transforms* defined in [4]. The regularized $\mathcal{MS}$ then captures salient structures of the shape in a simplified form. *(iii)* The $\mathcal{MS}$ can be further reduced into a succinct one-dimensional *graph* like structure by keeping the 2D sheet topology at the medial curves (and compress geometry information), which enables to adopt an efficient graph-based algorithm for matching. *(iv)* Finally, general practical data such as unorganized points or partial meshes can be handled without restrictions [4].

In this paper, we use the $\mathcal{MS}$ transitions and transforms as a main tool for shape matching, in addition to their use in $\mathcal{MS}$ regularization in [4]. In 2D, the $\mathcal{MA}$ transitions are exploited to navigate through the 'edits' of $\mathcal{MA}$ topologies (Fig.2) in exploring candidate deformation paths in matching two shapes [19]. We further extend this idea in 3D. Instead of matching the skeletal graph itself, we define a **dissimilarity** measure to reflect the *optimal* deformation cost between two shapes. The computational cost to explore all possible edits to match two $\mathcal{MS}$ hypergraphs is high [3, 19], we thus refer to a *sub-optimal* solution. We first regularize the $\mathcal{MS}$ hypergraphs to capture the *qualitative structure* of the shape and adopt a graph-matching scheme.

4. Graph Matching the $\mathcal{MS}$ Hypergraphs

Graphs are powerful data structures useful in matching and recognition [6]. The computational intractability of graph matching has led to the development of several classes of sub-optimal algorithms, including *search-oriented* methods using heuristics to explore the state space, and nonlinear optimization methods such as *relaxation labelling*. There are approaches using eigenvalue decomposition, neural networks, and linear programming, *etc*.

The *graduated assignment* ($\mathcal{GA}$) [13] is a relaxation-based energy-minimizing graph matching algorithm suitable to integrate with our approach. We adopt it to match the $\mathcal{MS}$'s for several reasons: *(i)* It enforces a *two-way* assignment via *"softassign"* [22], in contrast to *relaxation labelling* which enforces only an one-way assignment. It is extensible to a *three-way* assignment to match the hypergraphs, a strong fit in our case (§5.3). *(ii)* It avoids poor local minimum by a *graduated convexity* continuation technique [13]. *(iii)* It handles missing/extra nodes and links to stabilize the matching under noisy conditions, a factor essential in matching shapes. *(iv)* The formulation can be adapted to take into account shape deformations represented by continuous variables. *(v)* Finally, the computational time is comparable to other popular techniques.

4.1. A Review of the $\mathcal{GA}$ Graph Matching Algorithm

The basic setup of the $\mathcal{GA}$ [13] is to associate the nodes in two graphs G and $\overline{G}$ by a *match matrix* $\mathbf{M}$, where 1 represents a match of two nodes and 0 otherwise, Fig.3. $\mathbf{M}$ is a *permutation matrix* if the number of nodes in two graphs are equal. A *slack* row and column are added to $\mathbf{M}$ to represent missing/extra nodes. We refer to the nodes by G_a and $\overline{G}_i$, and the links by G_{ab} and $\overline{G}_{ij}$, respectively, where $a, b = 1, \ldots, A$, and $i, j = 1, \ldots, I$, *i.e.*,

$$\mathbf{M}_{ai} = \begin{cases} 1 & \text{if the node } a \in G \text{ corresponds to node } i \in \overline{G} \\ 0 & \text{otherwise.} \end{cases}$$

An objective *energy function* $E(\mathbf{M})$ is defined for each possible assignment $\mathbf{M}$. For a *quadratic* matching problem:

$$E(\mathbf{M}) = \sum_{i=1}^{I}\sum_{a=1}^{A}\sum_{b=1}^{A}\sum_{j=1}^{I} \mathbf{M}_{ai}\mathbf{M}_{bj}L_{aibj} \;, \tag{1}$$

where L_{aibj} represents the *compatibility* between links G_{ab} and $\overline{G}_{ij}$. Maximizing E yields the best matching between G and $\overline{G}$. A significant idea in [13] is to extend the above discrete assignment problem to a *continuous* one by embedding it into a large space, where *gradient descent* can be performed to iteratively move from one assignment to another. In this context, the continuous matching matrix $\mathbf{M}$ takes values between 0 and 1, satisfying the constraint of being a *doubly stochastic* matrix: $\sum_a \mathbf{M}_{ai} = 1$ and $\sum_i \mathbf{M}_{ai} = 1$ [22]. The $\mathcal{GA}$ algorithm then *differentially* moves from one assignment $\mathbf{M}$ to another, guided by refining the matching energy $E(\mathbf{M})$ in a *graduated non-convexity* setting [13], which slowly modifies $\mathbf{M}$ towards a 0 or 1 discretization. The Taylor expansion of E is:

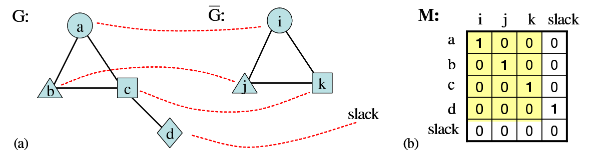

Figure 3. (a) The matching of two synthetic graphs G and $\bar{G}$, where the extra node is matched into the *slack* column of $\mathbf{M}$.

$$E(\mathbf{M}) = E(\mathbf{M}^0) + \sum_{a=1}^{A}\sum_{i=1}^{I} Q_{ai}(\mathbf{M}_{ai} - \mathbf{M}^0_{ai}), \tag{2}$$

where the derivative matrix Q is:

$$Q = \frac{\partial E}{\partial \mathbf{M}_{ai}}\bigg|_{\mathbf{M}=\mathbf{M}^0} = \sum_{b=1}^{A}\sum_{j=1}^{I} \mathbf{M}^0_{bj}\, L_{aibj}. \tag{3}$$

The problem of maximizing E is then turned into maximizing $\sum_{a=1}^{A}\sum_{i=1}^{I} Q_{ai}\,\mathbf{M}_{ai}$, which is (again) an *assignment problem* [13] solvable by softassign. The $\mathcal{GA}$ algorithm iteratively (and gradually) modifies an initial $\mathbf{M}$ toward discretization by decreasing a parameter T ('temperature' in annealing), which controls the convexity of the energy landscape to avoid poor local minima. In each iteration, $\mathbf{M}$ is best estimated and normalized toward a final assignment.

4.2. Extending $\mathcal{GA}$ to Match $\mathcal{MS}$ Hypergraphs

The original $\mathcal{GA}$ is shown to be successful in matching two *attributed relational graphs* by combining the energies of node-to-node assignments (1^{st}-order term) with link-to-link assignments (2^{nd}-order term) [13]. In our case to match the hypergraphs with yet another dimension, we introduce a 3^{rd}*-order* assignment to match the medial sheets (hyperlinks) of the $\mathcal{MS}$ hypergraph [5]. While the medial sheets could contain complex topology in general, we indirectly match them by matching individual *corners* of the sheet (where medial curves intersect), whose overall effects accumulate to match the sheets.

In defining the energy E to match the $\mathcal{MS}$ hypergraphs, ideally E should reflect the true similarity between two shapes. However it is difficult to model the exact shape variations in practice. Instead, we make E reflect the *component-wise* compatibility between two hypergraphs $\mathcal{MS}$ and $\overline{\mathcal{MS}}$, which composes of two measures to reflect the *structural* and *parametric* variations. We then optimize the overall similarity by summing up all compatibility measures (§5).

We further exploit two thoughts to improve the $\mathcal{MS}$ hypergraph matching scheme. First, a *square-root distance* $d^s = \sqrt{|m_1 - m_2|}$ is used for comparison between two measures m_1 and m_2, motivated by three reasons [20]: *(a)*

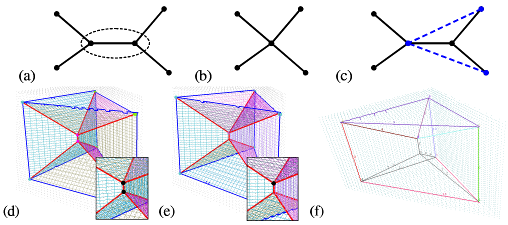

Figure 4. The adding of virtual links correlates a *contracting* edit of (a-c) a graph link in 2D and (d-e) a curve of a hypergraph in 3D. (f) shows the matching of two $\mathcal{MS}$ hypergraphs across a *contract transform* [4], where the mismatch links are shown in gray.

a re-interpretation of Weber's law, *(b)* to maintain sensitivity to variations when two items are close, and *(c)* to reduce sensitivity when two items are very distant. We found this norm outperforms the weighted distance ($\frac{|m_1-m_2|}{max(m_1,m_2)}$) and other metrics [20]. Second, a set of *virtual links* [20] (in addition to the medial curves/links) is introduced which improves the matching robustness in tolerating structural differences (detailed below).

The $\mathcal{MS}$ matching approach is described as follows. We first regularize the $\mathcal{MS}$ by applying a set of transforms [4] to simplify it across generic transitions, which bring the two hypergraphs close in structure. The remaining difference is then matched to evaluate similarity. Since the $\mathcal{GA}$ handles missing and extra nodes (by the slack variables), we can enforce the matching across structural differences by "simulating" them across remaining transitions, *e.g.*, by connecting the nodes that will be brought together in the future transitions by adding *virtual links*, Fig.4, which simulates the effect of applying further $\mathcal{MS}$ transforms [3]. We only add virtual links to the configurations that warrant high possibility of transitions (with slight shape variations). This helps to explore the better but difficult-to-compute *edit distance* [19] in the candidate spaces to improve the matching as well as its robustness. It essentially complements the one-to-one assignment nature of the $\mathcal{GA}$ graph-matching scheme in handling structural variations.

5. Matching the $\mathcal{MS}$ Hypergraphs: Componentwise Compatibility Measures

This section describes the component-wise node-to-node, curve-to-curve, and sheet-to-sheet compatibility measures to match two hypergraphs $\mathcal{MS}$ and $\overline{\mathcal{MS}}$. The compatibilities are defined for all assignments $\mathbf{M}$, such that the $\mathcal{GA}$ can migrate through the space to produce a final correspondence. The $\mathcal{MS}$ hypergraphs can be matched **structurally**, *e.g.*, whether a link/hyperlink exists or whether their types are consistent or not. They can also be matched **parametrically**, *i.e.*, by quantitatively comparing the attributes along the medial curves/sheets to define a metric. Consider that a shape can deform via a sequence of (canonical) transformations, dissimilarity between shapes is the amount of *minimum* changes in the space of transformations. We briefly consider the following measures: *(a) stretching* or *compressing*: the similarity is translated into the length comparison of medial curves; *(b) fattening* or *thinning*: the similarity is measured by the shock acceleration functions; *(c) bending* which affects the curvature along the $\mathcal{MA}$. These terms are measured in all hypergraph components (namely, all pairs of nodes, curves, and sheets) in order to compute the compatibility estimate. Table 1 overviews the main terms.

Table 1. Compatibility terms used in matching $\mathcal{MS}$ hypergraphs.

	Structural similarity	Parametric similarity
node	node type, incident curve types	radius r, ∇r along incident curves, angle of curves at sheet corners
curve	existence, curve type, ending node types, orientation	$\sum r_i$, edit distance (elastic deformations), Euclidean distance
sheet (corner)	existence, incident curve types, corner node type	corner angle, radius r, ∇r at the corner, $\sum r_i$ (to approximate volume)

5.1. The first-order node compatibility ($\mathcal{N}$)

Two shock nodes $N_a \in \mathcal{MS}$ and $\bar{N}_i \in \overline{\mathcal{MS}}$ are compared **structurally** on their node types and incident shock curve types. We observe two sets of shock nodes: one from the classification of the **5** *generic* $\mathcal{MA}$ nodes in [11] and the other of *high-order* nodes produced by applying $\mathcal{MS}$ transforms in the regularization process [4]. The above two sets can be re-organized into two categories, namely the $A_1^m A_3$ and A_1^n, where $m, \bar{m} \geq 1$ and $n, \bar{n} \geq 4$ [3]. We propose to penalize the difference in node types with normalization as:

$$d[A_1^m A_3, A_1^n] = \frac{max(|n-m|,3)}{max(n,m+3)}, d[A_1^n, A_1^{\bar{n}}] = \frac{|n-\bar{n}|}{max(n,\bar{n})},$$
$$d[A_1^m A_3, A_1^{\bar{m}} A_3] = \frac{|m-\bar{m}|}{max(m,\bar{m})+3}, \quad (4)$$

where $d[\cdot,\cdot]$ denotes the *node type difference*. Table 2 lists the values of d between a few high-order $\mathcal{MS}$ nodes frequently observed in practice. The structural node compatibility $\mathcal{N}_s$ is defined as the compliment of node type difference,

$$\mathcal{N}_s[N_a, \bar{N}_i] = 1 - d[N_a, \bar{N}_i], \quad (5)$$

such that $\mathcal{N}_s = 1$ if the shock types are identical, and $\mathcal{N}_s = 0$ if the shock types have nothing in common.

We consider three main terms between two shock nodes in the **parametric** compatibility measure: *(i)* difference of the shock node *radius* r. *(ii)* difference of the *gradient* of radius (∇r) along their incident shock curves. Since there exist numerous incident curves at a shock node (for example, there are four curves at an A_1^4 node), we simply take the maximum and minimum measures ($\nabla r^+, \nabla r^-$) to compute the difference. [1] *(iii)* the *angles* a of the incident sheet

[1] Two additional reasons motivate this design: *(a)* the measures in between the maximum and minimum are less salient and less robust to compare; and *(b)* the $A_1 A_3$ node has only two incident curves, thus selecting two distinct curves is most general.

Table 2. Difference $d[\cdot,\cdot]$ between the *general* and *high-order* $\mathcal{MS}$ nodes used to measure their structural compatibility.

	A_1A_3	A_1^4	$A_1^2A_3$	A_1^5	$A_1^3A_3$	A_1^6	$A_1^4A_3$	A_1^7	$A_1^5A_3$	A_1^8	...	$A_1^mA_3$	A_1^n
A_1A_3	0	3/4	1/5	4/5	2/6	5/6	3/7	6/7	4/8	7/8	...	$\frac{m-1}{m+3}$	$\frac{n-1}{n}$
A_1^4		0	3/5	1/5	3/6	2/6	3/7	3/7	4/8	4/8	...	$\frac{m-1}{m+3}$	$\frac{n-4}{n}$
$A_1^2A_3$			0	3/5	1/6	4/6	2/7	5/7	3/8	6/8	...	$\frac{m-2}{m+3}$	$\frac{n-2}{n}$
A_1^5				0	3/6	1/6	3/7	2/7	3/8	3/8	...	$\frac{m-2}{m+3}$	$\frac{n-5}{n}$
$A_1^3A_3$					0	3/6	1/7	4/7	2/8	5/8	...	$\frac{m-3}{m+3}$	$\frac{n-3}{n}$
A_1^6						0	3/7	1/7	3/8	2/8	...	$\frac{m-3}{m+3}$	$\frac{n-6}{n}$
$A_1^4A_3$							0	3/7	1/8	4/8	...	$\frac{m-4}{m+3}$	$\frac{n-4}{n}$
A_1^7								0	3/8	1/8	...	$\frac{m-4}{m+3}$	$\frac{n-7}{n}$
$A_1^5A_3$									0	3/8	...	$\frac{m-5}{m+3}$	$\frac{n-5}{n}$
A_1^8										0	...	$\frac{m-5}{m+3}$	$\frac{n-8}{n}$
...											0	...	...
$A_1^mA_3$												0	...
A_1^n													0

corner between a pair of shock curves. Again there exists numerous sheet corners at each node (for example, there are six sheets intersecting at an A_1^4 node), we only take the maximum and minimum measures (a^+, a^-). The parametric node compatibility $\mathcal{N}_p$ is then:

$$\mathcal{N}_p[N_a, \bar{N}_i] = 1 - w_r^n \cdot d^s[r(N_a), r(\bar{N}_i)] - \frac{w_g^n}{2} d^s[\nabla r^+(N_a), \nabla r^+(\bar{N}_i)] - \frac{w_g^n}{2} d^s[\nabla r^-(N_a), \nabla r^-(\bar{N}_i)] - \frac{w_a^n}{2} d^s[a^+(N_a), a^+(\bar{N}_i)] - \frac{w_a^n}{2} d^s[a^-(N_a), a^-(\bar{N}_i)]. \quad (6)$$

where the weighting constants $w_r^n = 0.5$, $w_g^n = 0.25$, $w_a^n = 0.25$ specify the relative importance between different measures. Finally, the 1^{st}-order node compatibility for hypergraph matching is:

$$\mathcal{N}_{ai}[N_a, \bar{N}_i] = \mathcal{N}_s \cdot \mathcal{N}_p. \quad (7)$$

To illustrate the performance of the above measure, we investigate in detail an example in matching two bone shapes in Fig.1 throughout this paper. A set of manually selected pairs of matching nodes is used as ground truth. The resulting node compatibility is shown in Table 3.

5.2. The second-order curve compatibility ($\mathcal{L}$)

Two shock curves $C_{ab} \in \mathcal{MS}$ and $\bar{C}_{ij} \in \overline{\mathcal{MS}}$ are compared **structurally** on their curve types, curve orientations, and end node types. If C_{ab} or $\bar{C}_{ij}$ is missing or the curve types are different, $\mathcal{L}_s[C_{ab}, \bar{C}_{ij}] = 0$; otherwise,

$$\mathcal{L}_s = 1 - w_n^l \cdot d[N_a, \bar{N}_i] - w_n^l \cdot d[N_b, \bar{N}_j], \quad (8)$$

where the weighting parameter $w_n^l = 0.2$ penalizes the difference between the ending shock node types.

We consider three main terms in the **parametric** compatibility measure between C_{ab} and $\bar{C}_{ij}$: *(i) integration of shock radius* along the curves to reflect corresponding shape volume,

$$V(C) = \int_{s \in C} r \cdot ds \approx \sum_{k=1}^{n_{sample}} r(C_k), \quad (9)$$

with a normalization done by dividing by a maximum value in $\mathcal{MS}$ and $\overline{\mathcal{MS}}$, such that $0 \le V^n(C_{ab}), V^n(\bar{C}_{ij}) \le 1$. *(ii) edit distance* $d_{ed}[C_{ab}, \bar{C}_{ij}]$ optimizing the *elastic* deformation between the two curves to obtain an optimal alignment minimizing the following energy terms via *dynamic programming* (D.P.) [18, 19]: *(a) stretching* difference $ds(C_p, C_{p\text{-}1}, \bar{C}_q, \bar{C}_{q\text{-}1})$, where p, q are indices of the fine-scale mesh vertices along the curves, *(b) bending* difference dt, *(c)* shock *radius* difference dr, and *(d)* the *angle* difference da between incident sheets, Fig.5a. Specifically, two shock curves are matched both as space curves (in their elastic bending and stretching terms) and as a joined skeleton-boundary matching to approximate their similarity. [2]

$$\min \ \left\{ [ds + dr] + w_a^l \cdot [dt + da] \right\}, \quad (10)$$

where the weighting constant $w_a^l = 0.1$ specifies the ratio between the *length* and *angle* measures (the value is relative to the *scale* of shapes as in [19]). *(iii)* an optional *Euclidean distance* d_{Eu} available from the correspondence of the above edit-distance matching, Fig.5b. The parametric curve compatibility is then:

$$\mathcal{L}_p[C_{ab}, \bar{C}_{ij}] = 1 - w_d^l \cdot d_{ed}[C_{ab}, \bar{C}_{ij}] - w_e^l \cdot d_{Eu}[C_{ab}, \bar{C}_{ij}] - w_v^l \cdot \left|V(C_{ab}) - V(\bar{C}_{ij})\right|, \quad (11)$$

where the weighting constants are chosen as $w_d^l = 0.33$, $w_e^l = 0.33$, and $w_v^l = 0.34$. Finally, the 2^{nd}-order link (curve) compatibility for hypergraph matching is:

$$\mathcal{L}_{aibj}[C_{ab}, \bar{C}_{ij}] = \mathcal{L}_s \cdot \mathcal{L}_p. \quad (12)$$

Table 4 demonstrates high performance of the above compatibility measure in matching shock curves. Further experiments (not shown) indicated that it provides a strong clue to improve the proposed matching performance.

Table 3. An example *node compatibility* table ($\mathcal{N}_{ai}$) in matching the $\mathcal{MS}$ nodes in Fig.1 (value in 1/100). Ground truth pairs with high compatibilities are highlighted in parentheses (boldface in blue). A few *erroneous* matches with value higher than the ground true are underlined in red.

	n51	n59	n69	n76	n80	n83	n89	n90	n91	n92	n94
N00	44	41	41	55	49	53	24	14	16	15	35
N09	<u>57</u>	<u>68</u>	<u>67</u>	<u>68</u>	<u>73</u>	<u>56</u>	**(37)**	10	12	11	22
N13	68	68	<u>77</u>	54	**(72)**	57	40	11	12	12	28
N19	54	62	56	**(69)**	68	49	39	09	10	10	22
N25	59	67	73	45	63	50	34	10	11	10	23
N37	55	58	59	72	66	52	39	09	10	10	23
N46	12	10	13	11	12	12	0	75	69	71	15
N48	14	12	11	12	11	13	0	71	**(86)**	80	16
N51	20	19	18	20	20	20	37	0	0	0	22
N52	<u>20</u>	<u>18</u>	**(17)**	<u>22</u>	<u>19</u>	<u>18</u>	<u>34</u>	0	0	0	<u>19</u>
N56	19	14	15	16	14	18	0	44	53	55	14
N57	12	10	10	11	10	11	0	**(80)**	73	73	14

5.3. The third-order sheet (corner) compatibility ($\mathcal{C}$)

Two shock sheet corners $S_{abc} \in \mathcal{MS}$ and $\bar{S}_{ijk} \in \overline{\mathcal{MS}}$ are compared **structurally** on their boundary shock *curve type* and the shock *node types* at the corner. Specifically, if S_{abc} or $\bar{S}_{ijk}$ is missing or $(C_{ab}, \bar{C}_{ij})$ or $(C_{bc}, \bar{C}_{jk})$ are of different types, $\mathcal{C}_s = 0$; otherwise,

[2] We use an *alignment curve* α to represent the result of the correspondence between C_{ab} and $\bar{C}_{ij}$ [19, 18]. Details on computing the differential terms of the shock curves as space curves (stretching and derivatives ϕ, ϕ_s, ϕ_{ss}, bending and derivatives θ, θ_s, θ_{ss}, tangent T, normal N, binomal B, curvature κ, and torsion τ) are omitted due to space limit.

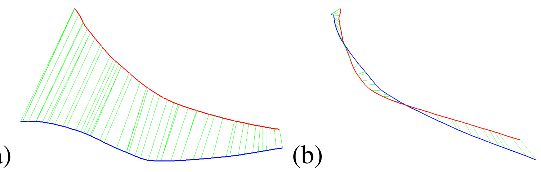

(a) (b)

Figure 5. (a) The D.P. matching of two shock curves by considering their geometry as 3D space curves together with other shock attributes. (b) The rigid alignment by using the D.P. correspondence to compute the optimal Euclidean distance between them.

Table 4. An example *curve compatibility* table ($\mathcal{L}_{aibj}$) in matching the $\mathcal{MS}$ curves in Fig.1 (value in $1/100$). Ground truth pairs are highlighted in parentheses (boldface in blue). Only four *erroneous* matches are underlined in red, suggesting a strong discriminative performance.

	c00	c25	c47	c72	c74	c84	c93	c94	c95	c99	c100	c101	c102	c103	c105	c107
C01	00	00	00	00	60	27	19	00	78	39	38	40	59	79	52	68
C02	11	68	04	84	00	00	00	63	00	00	00	00	00	00	00	00
C10	81	24	39	19	00	00	00	06	00	00	00	00	00	00	00	00
C14	10	61	03	80	00	00	00	55	00	00	00	00	00	00	00	00
C26	15	80	09	61	00	00	00	27	00	00	00	00	00	00	00	00
C33	66	45	61	37	00	00	00	08	00	00	00	00	00	00	00	00
C39	00	00	00	00	51	23	15	00	75	40	41	42	46	54	50	48
C50	00	00	00	00	**(91)**	32	19	00	75	44	50	44	69	51	35	43
C59	00	00	00	00	53	24	17	00	68	51	33	50	73	96	58	58
C60	00	00	00	00	48	23	15	00	61	48	31	49	71	92	63	61
C64	00	00	00	00	23	53	65	00	20	57	33	55	35	25	18	13
C65	00	00	00	00	73	39	21	00	51	44	74	42	45	36	24	26
C66	00	00	00	00	51	23	15	00	67	35	36	36	53	73	48	69
C67	00	00	00	00	51	23	15	00	67	35	36	36	53	73	48	69
C68	00	00	00	00	54	**(39)**	23	00	40	38	72	36	35	29	19	20
C69	00	00	00	00	40	37	24	00	32	63	47	58	51	39	18	20
C70	00	00	00	00	43	43	24	00	31	65	48	56	50	39	18	19
C71	00	00	00	00	59	72	41	00	41	36	**(69)**	32	34	26	14	18

$$\mathcal{C}_s(S_{abc},\bar{S}_{ijk}) = 1 - w^c_{n1} d[N_b,\bar{N}_j] - w^c_{n2} d[N_a,\bar{N}_i] - w^c_{n2} d[N_c,\bar{N}_k]$$

with weighting constants $w^c_{n1} = 0.8$ and $w^c_{n2} = 0.1$. We make $w^c_{n1} + 2w^c_{n2} = 1$ to completely determine the structural compatibility, due to the fact that there is no structural information from the medial sheets (all of A_1^2 type).

We consider four main terms in the **parametric** compatibility measure: *(i)* shock radius r of the corner node (N_b and $\bar{N}_j$), *(ii)* gradient of shock radius (∇r) along the sheet (bisector curve at the corner), *(iii)* corner angle between $\angle(S_{abc})$ and $\angle(\bar{S}_{ijk})$, *(iv)* corresponding shape volume V of the sheet (integration of shock radius). The parametric sheet (corner) compatibility $\mathcal{C}_p$ is defined as:

$$\begin{aligned}\mathcal{C}_p(S_{abc},\bar{S}_{ijk}) &= 1 - w^c_r d^s[r(N_b), r(\bar{N}_j)] - w^c_g d^s[\nabla r(S), \nabla r(\bar{S})] \\ &- w^c_a d^s[\angle(S_{abc}), \angle(\bar{S}_{ijk})] - w^c_v d^s[V(S), V(\bar{S})]. \qquad (13)\end{aligned}$$

where $w^c_r = 0.2$, $w^c_g = 0.3$, $w^c_a = 0.3$, and $w^c_v = 0.2$. Finally, the third-order sheet (corner) compatibility is:

$$\mathcal{C}_{aibjck}[S_{abc},\bar{S}_{ijk}] = \mathcal{C}_s \cdot \mathcal{C}_p. \qquad (14)$$

Table 5 examines the performance of this design. Further experiments (not shown) suggested that $\mathcal{C}_{aibjck}$ provides strong clues to boost the hypergraph matching.

6. Computing the Post-Matching Similarity

In this section, we derive a *global* similarity measure as an energy functional using the matching result and the compatibilities defined in §5. The energy $E(\mathbf{M})$ between two

Table 5. An example *sheet (corner) compatibility* table ($\mathcal{C}_{aibjck}$) in matching the $\mathcal{MS}$ sheets in Fig.1 (values in $1/100$). Ground truth pairs are highlighted in parentheses (boldface in blue). A few *erroneous* matches are underlined in red.

	c76	c91	c91	c91	c89	c91	c91	c89	c90	c90	c90	c51	c90	c90	c89	c91	c92	c92	c94	c59	c69	c80	c94	c92	c90	c83
C[00-46-48]	0	48	53	59	0	52	49	0	47	55	53	0	54	56	0	52	45	53	0	0	0	0	15	57	60	0
C[00-46-56]	0	40	47	46	0	37	49	0	39	42	46	0	51	48	0	52	44	56	0	0	0	0	15	41	53	0
C[00-46-57]	0	52	63	67	0	48	53	0	50	63	56	0	59	62	0	61	40	55	0	0	0	0	18	58	70	0
C[00-51-56]	21	0	0	0	13	0	0	17	0	0	0	22	0	0	20	0	0	0	27	16	19	18	0	0	0	0
C[09-52-57]	28	0	0	0	25	0	0	31	0	0	0	20	0	0	34	0	0	0	33	21	23	28	0	0	0	0
C[09-57-19]	0	68	74	63	0	70	63	0	73	66	62	0	62	68	0	62	41	52	0	0	0	0	16	46	46	0
C[09-57-46]	0	54	62	66	0	57	64	0	58	69	67	0	64	64	0	70	38	49	0	0	0	0	22	43	57	0
C[09-57-52]	0	54	62	60	0	59	68	0	60	63	69	0	64	59	0	67	36	49	0	0	0	0	16	43	46	0
C[13-25-52]	41	0	0	0	26	0	0	29	0	0	0	55	0	0	26	0	0	0	21	48	53	43	0	0	0	0
C[13-48-37]	0	68	69	57	0	**(74)**	58	0	75	61	58	0	55	60	0	58	46	52	0	0	0	0	14	43	41	0
C[13-48-46]	0	53	60	65	0	57	66	0	59	71	69	0	64	63	0	**(73)**	40	52	0	0	0	0	22	43	60	0
C[13-48-56]	0	56	**65**	63	0	61	71	0	63	68	71	0	69	65	0	69	44	47	0	0	0	0	20	39	50	0
C[19-51-56]	27	0	0	0	25	0	0	25	0	0	0	26	0	0	29	0	0	0	23	26	25	26	0	0	0	0
C[19-57-46]	0	53	60	64	0	56	73	0	58	68	74	0	**(70)**	64	0	73	42	49	0	0	0	0	22	41	50	0
C[19-57-52]	0	63	64	55	0	74	55	0	**(66)**	58	55	0	54	57	0	56	46	53	0	0	0	0	13	46	40	0
C[25-13-48]	50	0	0	0	30	0	0	34	0	0	0	73	0	0	29	0	0	0	28	59	67	52	0	0	0	0
C[25-52-37]	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	13
C[25-52-57]	0	7	7	10	0	8	12	0	8	12	12	0	11	10	0	13	8	5	0	0	0	0	29	2	5	0
C[37-48-46]	0	54	63	61	0	57	**(70)**	0	58	62	67	0	67	64	0	71	44	46	0	0	0	0	22	38	47	0
C[37-48-56]	0	**(61)**	58	57	0	69	55	0	59	58	55	0	56	60	0	56	50	52	0	0	0	0	19	43	41	0
C[37-52-57]	28	0	0	0	30	0	0	34	0	0	0	20	0	0	36	0	0	0	24	21	21	25	0	0	0	0
C[46-00-51]	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	56
C[46-48-56]	0	43	49	**54**	0	40	52	0	41	51	56	0	59	52	0	51	36	44	0	0	0	0	14	61	68	0
C[46-56-48]	0	33	37	40	0	35	33	0	32	37	35	0	37	39	0	38	43	41	0	0	0	0	11	57	56	0
C[46-56-51]	0	38	36	30	0	32	31	0	31	28	30	0	33	31	0	34	35	45	0	0	0	0	8	45	42	0
C[46-57-52]	0	43	51	50	0	40	41	0	42	**(47)**	43	0	45	46	0	46	33	51	0	0	0	0	9	68	63	0
C[48-37-52]	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	47
C[48-46-56]	0	42	44	48	0	40	42	0	42	45	42	0	48	51	0	39	42	60	0	0	0	0	11	56	60	0
C[48-46-57]	0	33	39	42	0	32	39	0	31	39	40	0	42	40	0	48	31	38	0	0	0	0	15	49	64	0
C[48-56-51]	0	43	50	44	0	41	47	0	42	41	44	0	48	46	0	49	34	42	0	0	0	0	9	54	59	0
C[51-19-57]	**(62)**	0	0	0	27	0	0	28	0	0	0	47	0	0	30	0	0	0	28	48	45	56	0	0	0	0
C[52-09-57]	58	0	0	0	22	0	0	26	0	0	0	52	0	0	30	0	0	0	36	53	57	75	0	0	0	0
C[56-46-57]	0	52	56	64	0	55	50	0	48	57	53	0	57	60	0	53	49	61	0	0	0	0	15	57	54	0

hypergraphs $\mathcal{MS}$ and $\overline{\mathcal{MS}}$ is defined as:

$$E[\mathcal{MS},\overline{\mathcal{MS}},\mathbf{M}] = w^N \cdot E_N + w^L \cdot E_L + w^C \cdot E_C. \qquad (15)$$

where $w^N = 0.3$, $w^L = 0.5$, $w^C = 0.4$ are the weighting constants between the node (E_N), curve (E_L), and sheet corner (E_C) energies:

$$E_N = \sum_{a=1}^{A}\sum_{i=1}^{I} \mathbf{M}_{ai}\mathcal{N}_{ai}[N_a,\bar{N}_i],$$

$$E_L = \sum_{a=1}^{A}\sum_{i=1}^{I}\sum_{b=1}^{A}\sum_{j=1}^{I} \mathbf{M}_{ai}\mathbf{M}_{bj}\mathcal{L}_{aibj}[C_{ab},\bar{C}_{ij}],$$

$$E_C = \sum_{a=1}^{A}\sum_{i=1}^{I}\sum_{b=1}^{A}\sum_{j=1}^{I}\sum_{c=1}^{A}\sum_{k=1}^{I} \mathbf{M}_{ai}\mathbf{M}_{bj}\mathbf{M}_{ck}\mathcal{C}_{aibjck}[S_{abc},\bar{S}_{ijk}].$$

Our initial implementation of Eq. (15) reveals that the energy defined in this way varies with the number of nodes and links in each hypergraph, thus the similarity between pairs of shapes could not be universally compared. This motivates a *normalized similarity energy* E^n [20]:

$$\begin{aligned} E^n[\mathcal{MS},\overline{\mathcal{MS}},\mathbf{M}] = w^N \sum_{a=1}^{A}\sum_{i=1}^{I} \frac{\mathbf{M}_{ai}\mathcal{N}_{ai}}{\mathbf{M}_{ai}} + \qquad (16) \\ w^L \sum_{a=1}^{A}\sum_{i=1}^{I}\sum_{b=1}^{A}\sum_{j=1}^{I} \frac{\mathbf{M}_{ai}\mathbf{M}_{bj}\mathcal{L}_{aibj}}{\mathbf{M}_{ai}\mathbf{M}_{bj}\mathcal{L}^{max}_{aibj}} + \\ w^C \sum_{a=1}^{A}\sum_{i=1}^{I}\sum_{b=1}^{A}\sum_{j=1}^{I}\sum_{c=1}^{A}\sum_{k=1}^{I} \frac{\mathbf{M}_{ai}\mathbf{M}_{bj}\cdot\mathbf{M}_{ck}\mathcal{C}_{aibjck}}{\mathbf{M}_{ai}\mathbf{M}_{bj}\cdot\mathbf{M}_{ck}\mathcal{C}^{max}_{aibjck}}, \end{aligned}$$

where $\mathcal{L}^{max}_{aibj} = 0$ if C_{ab} or $\bar{C}_{ij}$ is missing, or 1 otherwise; $\mathcal{C}^{max}_{aibjck} = 0$ if $C_{ab}, C_{bc}, \bar{C}_{ij}, \bar{C}_{jk}, S_{abc}$ or $\bar{S}_{ijk}$ is missing, or 1 otherwise. Mis-matches in the slacks of $\mathbf{M}$ are penalized with compatibility 0. This new design ensures $0 \leq E^n \leq 1$. We use E in the $\mathcal{GA}$ matching process and use E^n for the final similarity computation.

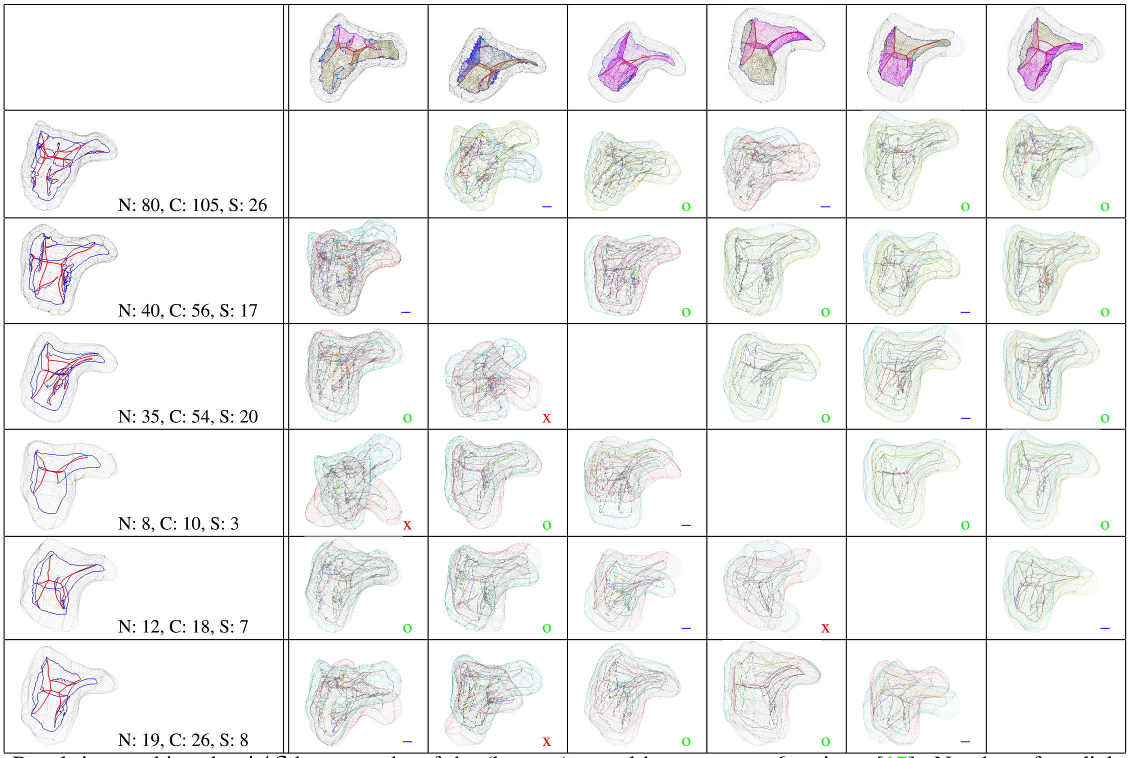

Figure 6. Result in matching the $\mathcal{MS}$ hypergraphs of the (hamate) carpal bones across 6 patients [17]. Number of medial nodes (N), curves (C), and sheets (S) of the $\mathcal{MS}$'s after regularization are labelled on the left. We use the node correspondence to compute a rigid transformation to align the bones for visualization. Closely aligned bones are manually marked with 'o' (16 ones), while roughly aligned ones are marked with '–' (10 ones) and misaligned cases are marked with 'x' (4 ones), out of a visual examination. Since visual misalignment does not necessarily indicate a wrong match, this experiment only evaluates qualitatively how the method performs on matching similar shapes with large skeletal structure variations.

7. Results and Discussions

We perform an initial experiment on a human wrist bone dataset (courtesy of Dr. Crisco, RI Hospital [17]) to examine the shape variations of the carpal bones across patients. In Fig.6, the left-hand *hamates* of 6 females are matched against each other. Note that the graph matching in this case is challenging. The regularized $\mathcal{MS}$ hypergraphs contain large topological variations: One bone contains only 8 node/10 curves/3 sheets, while another contains 80 nodes/105 curves/26 sheets. The proposed $\mathcal{GA}$ matching maps most mismatches into the slacks, leading to many correct matches. We observe *asymmetry* matching results due to the rounding of $\mathbf{M}$ into discrete assignments.

We have also tested our matching approach on a small database (Fig.7) composed of 20 shapes in 5 categories. The artificial shapes are chosen to examine the matching across particular $\mathcal{MS}$ transitions in isolation [4], while one hand model is collected online (from Polhemus) and the others are generated from laser scans. Note that similarity measures of within-category shapes are much higher than non-category shapes, indicating the potential in applications such as shape retrieval and recognition.

The matching of two scaffolds typically takes 5 to 10 minutes to run on a computer with moderate configuration. The introduction of the 3^{rd}-order energy makes the complexity of Eq. (15) to be $O(A^3 I^3)$. However this does not affect the computation time too much due to an efficient encoding of sparsity (Observe the sparsity in Table 5).

Conclusion: We have presented a medial axis based graph matching approach to measure 3D shape similarity by matching their $\mathcal{MS}$ hypergraphs. This is the first attempt exploiting the $\mathcal{MA}$ transitions [12] integrated with a hypergraph matching framework. We aim to include standard databases in our experiments and compare the result with other skeletal matching methods (*e.g.*, the curve skeleton [7] and medial surfaces [21]) in the future.

References

[1] S. Biasotti, D. Giorgi, M. Spagnuolo, and B. Falcidieno. Reeb graphs for shape analysis and applications. *Theoretical Computer Science*, 392(1-3):5–22, 2008. 2

[2] B. Bustos, D. A. Keim, D. Saupe, T. Schreck, and D. V. Vranić. Feature-based similarity search in 3D object databases. *ACM Comput. Survey*, 37(4):345–387, 2005. 2

[3] M.-C. Chang. *The Medial Scaffold for 3D Shape Modeling and Recognition*. Ph.D. dissertation, Division of Engineering, Brown University, Providence RI, 2009. 1, 3, 4

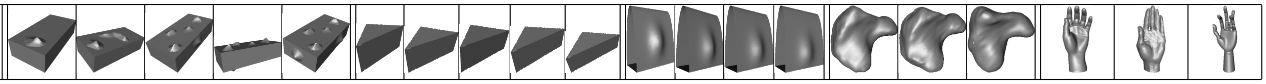

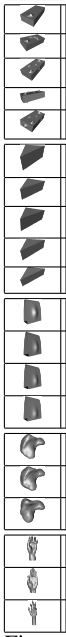

1.000	**0.870**	**0.710**	**0.692**	0.638	0.535	0.539	0.530	0.479	0.523	0.223	0.202	0.251	0.282	0.385	0.336	0.237	0.163	0.076	0.162
0.790	**1.000**	**0.700**	**0.660**	**0.647**	0.545	0.559	0.529	0.519	0.547	0.202	0.206	0.208	0.279	0.460	0.363	0.238	0.234	0.098	0.219
0.701	**0.660**	**1.000**	**0.869**	**0.955**	0.612	0.553	0.598	0.555	0.600	0.238	0.082	0.256	0.279	0.444	0.350	0.250	0.217	0.082	0.171
0.719	**0.699**	**0.843**	**1.000**	**0.877**	0.484	0.475	0.478	0.589	0.531	0.223	0.225	0.230	0.264	0.348	0.422	0.223	0.188	0.106	0.154
0.646	**0.661**	**0.953**	**0.893**	**1.000**	0.591	0.604	0.608	0.565	0.594	0.243	0.138	0.250	0.275	0.421	0.365	0.292	0.197	0.103	0.150
0.578	0.577	0.615	0.671	0.646	**1.000**	**0.917**	**0.934**	**0.862**	**0.730**	0.148	0.108	0.140	0.265	0.247	0.244	0.271	0.147	0.010	0.119
0.583	0.577	0.624	0.689	0.657	**0.908**	**1.000**	**0.924**	**0.899**	**0.823**	0.138	0.199	0.180	0.238	0.263	0.262	0.268	0.147	0.017	0.119
0.575	0.562	0.619	0.681	0.645	**0.925**	**0.918**	**1.000**	**0.885**	**0.729**	0.146	0.178	0.153	0.238	0.256	0.251	0.275	0.141	0.009	0.117
0.556	0.538	0.594	0.673	0.610	**0.867**	**0.907**	**0.896**	**1.000**	**0.854**	0.337	0.200	0.300	0.351	0.343	0.362	0.312	0.176	0.097	0.183
0.585	0.597	0.645	0.668	0.659	**0.734**	**0.832**	**0.733**	**0.818**	**1.000**	0.311	0.261	0.290	0.324	0.404	0.383	0.329	0.260	0.083	0.176
0.157	0.131	0.169	0.165	0.172	0.133	0.140	0.132	0.327	0.267	**1.000**	**0.613**	**0.903**	**0.505**	0.220	0.311	0.190	0.012	0.005	0.041
0.209	0.562	0.085	0.187	0.136	0.277	0.287	0.279	0.160	0.340	**0.662**	**1.000**	**0.597**	**0.496**	0.179	0.282	0.143	0.038	0.007	0.061
0.159	0.173	0.188	0.192	0.181	0.146	0.142	0.138	0.334	0.232	**0.931**	**0.596**	**1.000**	**0.527**	0.208	0.338	0.167	0.013	0.011	0.040
0.278	0.246	0.266	0.277	0.274	0.278	0.303	0.319	0.370	0.299	**0.492**	0.192	**0.493**	**1.000**	0.236	0.292	0.203	0.036	0.022	0.121
0.427	0.444	0.420	0.392	0.360	0.243	0.263	0.234	0.292	0.328	0.204	0.016	0.205	0.218	**1.000**	**0.483**	**0.458**	0.118	0.033	0.145
0.399	0.413	0.245	0.437	0.235	0.306	0.306	0.304	0.370	0.370	0.423	0.283	0.401	0.226	**0.513**	**1.000**	**0.493**	0.125	0.048	0.170
0.277	0.220	0.341	0.223	0.331	0.241	0.281	0.219	0.296	0.382	0.182	0.138	0.173	0.191	**0.528**	**0.493**	**1.000**	0.018	0.012	0.055
0.191	0.194	0.196	0.183	0.182	0.065	0.130	0.127	0.142	0.231	0.016	0.046	0.016	0.052	0.100	0.123	0.033	**1.000**	**0.287**	**0.311**
0.093	0.046	0.087	0.102	0.099	0.008	0.011	0.023	0.019	0.034	0.008	0.012	0.021	0.017	0.032	0.050	0.015	**0.274**	**1.000**	**0.239**
0.016	0.017	0.015	0.017	0.015	0.012	0.013	0.015	0.013	0.014	0.011	0.012	0.011	0.012	0.114	0.137	0.079	0.217	**0.258**	**1.000**

Figure 7. Similarity metric computed for a small shape database. Correct matches are highlighted in boldface in **black** (same object) and **blue**, while erroneous ones are in red. Observe that the *within-category* similarity is well-preserved and distinct from non-category shapes.

[4] M.-C. Chang and B. Kimia. Regularizing 3D medial axis using medial scaffold transforms. In *Proceedings of IEEE CVPR*, pages 1–8, 2008. 1, 2, 4, 7

[5] M.-C. Chang, F. Leymarie, and B. Kimia. 3D shape registration using regularized medial scaffolds. In *Proceedings of IEEE 3DPVT*, pages 987–994, 2004. 1, 3

[6] D. Conte, P. Foggia, C. Sansone, and M. Vento. Thirty years of graph matching in pattern recognition. *IJPRAI*, 18(3):265–298, 2004. 3

[7] N. D. Cornea, M. F. Demirci, D. Silver, A. Shokoufandeh, S. J. Dickinson, and P. B. Kantor. 3D object retrieval using many-to-many matching of curve skeletons. In *SMI*, pages 366–371, 2005. 2, 7

[8] T. Funkhouser, P. Min, M. Kazhdan, J. Chen, A. Halderman, D. Dobkin, and D. Jacobs. A search engine for 3D models. *ACM Transactions on Graphics*, 22(1):83–105, 2003. 1

[9] R. Gal and D. Cohen-Or. Salient geometric features for partial shape matching and similarity. *ACM Transactions on Graphics*, 25(1):130–150, 2006. 2

[10] P. Giblin and B. Kimia. On the intrinsic reconstruction of shape from its symmetries. *PAMI*, 25(7):895–911, 2003. 1

[11] P. Giblin and B. Kimia. A formal classification of 3D medial axis points and their local geometry. *PAMI*, 26(2):238–251, 2004. 1, 2, 4

[12] P. Giblin and B. Kimia. Transitions of the 3D medial axis under a one-parameter family of deformations. *PAMI*, 31(5):900–918, 2009. 1, 2, 7

[13] S. Gold and A. Rangarajan. A graduated assignment algorithm for graph matching. *PAMI*, 18(4):377–388, 1996. 1, 3

[14] C. Y. Ip and S. K. Gupta. Retrieving matching CAD models by using partial 3D pt. clouds. *CADA*, 4:629–638, 2007. 2

[15] A. E. Johnson and M. Hebert. Using spin images for efficient object recognition in cluttered 3D scenes. *PAMI*, 21(5):433–449, 1999. 2

[16] F. Leymarie and B. Kimia. The medial scaffold of 3D unorganized point clouds. *PAMI*, 29(2):313–330, 2007. 1, 2

[17] D. C. Moore, J. J. Crisco, T. G. Trafton, and E. L. Leventhal. A digital database of wrist bone anatomy and carpal kinematics. *J. Biomechanics*, 40(11):2537–42, 2007. 1, 7

[18] T. Sebastian, P. Klein, and B. Kimia. On aligning curves. *PAMI*, 25(1):116–125, January 2003. 5

[19] T. Sebastian, P. Klein, and B. Kimia. Recognition of shapes by editing their shock graphs. *PAMI*, 26:551–571, 2004. 1, 2, 3, 4, 5

[20] D. Sharvit, J. Chan, H. Tek, and B. B. Kimia. Symmetry-based indexing of image databases. *JVCIR*, 9(4):366–380, 1998. 2, 3, 4, 6

[21] K. Siddiqi, J. Zhang, D. Macrini, A. Shokoufandeh, S. Bouix, and S. Dickinson. Retrieving articulated 3-D models using medial surfaces. *Machine Vision and Applications*, 19:261–275, 2008. 1, 2, 7

[22] R. Sinkhorn. A relationship between arbitrary positive matrices and doubly stochastic matrices. *Annals of Mathematical Statistics*, 35:876–879, 1964. 3

[23] M. Styner, G. Gerig, S. Joshi, and S. Pizer. Automatic and robust computation of 3D medial models incorporating object variability. *IJCV*, 55(2-3):107–122, 2003. 2

[24] H. Sundar, D. Silver, N. Gagvani, and S. Dickinson. Skeleton based shape matching and retrieval. In *SMI*, pages 130–139, 2003. 2