

Knowledge-based Part Correspondence

Boaz J. Super, Ph.D.¹

Computer Science Department, University of Illinois at Chicago, Chicago, IL 60607 USA

Abstract

This paper presents a direct method for finding corresponding pairs of parts between two shapes. Statistical knowledge about a large number of parts from many different objects is used to find a part correspondence between two previously unseen input shapes. No class membership information is required. The knowledge-based approach is shown to produce significantly better results than a classical metric distance approach. The potential role of part correspondence as a complement to geometric and structural comparisons is discussed.

Keywords: part correspondence, shape matching, chance probability functions, nonaccidentalness, knowledge-based matching

¹ Present address: Motorola Labs, Schaumburg, IL 60196 USA. Email address: super1@motorola.com

1. Introduction

Recent research on shape matching has emphasized two shape matching problems. In *geometric shape comparison*, correspondences are found between points or other low-level spatial data of the shapes. The central challenge of correspondence matching is that low-level spatial data are relatively indistinct; therefore, many solution methods rely on global optimization [3,5] or point-centered contextual information [2,12] to disambiguate correspondences. In *structural shape comparison*, shape parts and their relationships are matched. The central challenge of structural matching is that structural representations are unstable (small geometric changes can cause large structural changes); therefore, solution methods typically perform approximate graph, tree, or sequence matching [9,10,21].

In this paper we explore two ideas. First, we propose a third shape matching problem, *part correspondence*, that is complementary to both geometric and structural matching. A part correspondence is a mapping between two sets of parts. Compared with points and other low-level geometric data, complex parts are less ambiguous and more salient perceptually. They also lend themselves naturally to handling occluded and articulated shapes, and to indexing object databases by parts [23]. Furthermore, in contrast with structural matching [1,8,9, 10,21,29], part correspondence does not match relationships, just parts, and accurate part correspondences can be obtained by the algorithm presented here. Part correspondence can be viewed as a type of shape comparison intermediate between geometric and structural comparisons.

The second idea explored here is that the widely-studied problem of matching shapes in isolation – that is, without considering knowledge about all shapes – is an artificial problem. This view is motivated by current understanding that biological vision systems exploit knowledge in performing visual processing tasks. In this paper, the matching of two shapes will be treated as a problem to be solved in the presence of a large base of stored knowledge. A specific contribution of this paper is the use of knowledge of a large set of shape parts, from many different objects, to define a probabilistic measure of part difference. A greedy algorithm for part correspondence is used. The central empirical demonstration of this paper is that when a

geometric measure of part difference is used, the algorithm fails; when the corresponding probabilistic measure of part difference is used, the algorithm succeeds. The large difference in outcome is due entirely to the use of different part difference measures; everything else is held constant. We suspect that direct part correspondence has been little used in computer vision because of the difficulty of correctly pairing parts of two shapes based only on the geometric differences of those parts.

No knowledge of shape classes is required by the method. That is, although shape examples are required, it is not necessary that the examples be labeled or classified, which saves the cost of annotating the training set. The method also does not keep track of which shapes the example parts come from. The knowledge consists entirely of the relationships of parts in a representation space instead of semantic knowledge of any kind.

To prevent confusion, note that the statistical model presented in this paper is not related to the class of methods known as "statistical shape modeling" (SSM) [6]. Unlike SSM, the current approach: (1) uses a database-wide model instead of a separate model for each class; (2) does not require accurate point correspondences; and (3) does not use class membership information.

The present method was developed as a bootstrapping tool to use in a larger, ongoing project on automatically constructing SSMs (class models) from examples. It was designed for the problem of matching shapes represented as contours, and not for the problems of matching shapes to images [6,8] or images to images. The current method works on sets of contour segments, which can come from shapes consisting of multiple open or closed contours. However, only single-contour shapes are supported by our current implementation. For recent work on parts-based matching of shapes under conditions of partial data deletion and clutter, see [22] (see also [11] which classifies several types of partial contour deletion tests used in the prior literature).

Shape comparison is not a single problem, but a collection of problems representing different ways to compare shapes. One goal of this paper is to suggest that part correspondence can be a valuable complement to point correspondence and structural matching. For example, part

correspondence can be used to divide a single whole-shape point correspondence problem into several easier partial-shape point correspondence problems. This will be demonstrated in Section 3. Part correspondences could also be used to constrain the set of candidate structural matches between two shapes. Finally, we hope that the examples presented in this paper will encourage thinking about using part correspondences directly in applications that do not require point or relational correspondences.

The statistical model used here is embodied in a set of *chance probability functions* (CPFs) [24]. The term "chance" in this context refers to a chance hypothesis, analogous to the chance or null hypothesis in nonaccidentalness [18,28]. We introduced CPFs based on whole planar shapes in previous research on shape retrieval and classification [24,27]. There, we demonstrated that the use of CPFs improved accuracy on the widely used MPEG-7 benchmark test of retrieval accuracy [16]. However, the earlier work was applicable only to matching whole, unoccluded shapes, and gave no information about the internal structure of the shapes. In this paper we use part-based CPFs instead of shape-based CPFs, and extend the knowledge-based matching idea to the new problem of finding part correspondences between two shapes.

It is interesting to note another recent approach to parts-based shape analysis [17], which matches two shapes or shape parts by simplifying them to increase their similarity. That method is based on earlier work which used structural matching of sequences of parts [15]. A different strategy is presented here: knowledge about many parts is used to compute the significances of measured part similarity values, and these significance values are used to more accurately correspond parts between two shapes. These two approaches, part simplification and part match significance, are not exclusive and can potentially be combined.

This paper is an extended version of a preliminary conference paper [26]. The new contributions include: (1) a formal mathematical description of part correspondence; (2) a formal description of the part correspondence algorithm; and (3) a new experiment demonstrating the combination of part correspondence with point correspondence. In addition the exposition has been expanded with new material in several sections. An earlier version of this paper is available

as a technical report [25].

2. Part Correspondence Method

This section describes the method. Section 2.1 summarizes the statistical model; Section 2.2 describes the generation and representation of parts; and Section 2.3 describes the use of the model to find part correspondences.

2.1 Chance probability functions

This section briefly summarizes the statistical model that was presented in detail in [24,27]. The basic intuition is that the distribution of shape parts in a high-dimensional part representation space is not uniform. Some shape parts are distinctive, appearing in one or a few object classes. This is more likely when the parts are complex. Other parts are relatively generic, appearing in many object classes. This is more likely when the parts are small, simple, or both. Because of the nonuniformity of the part distribution, a given metric distance between a pair of parts has different significance in high- and low-density regions of part space. In a high density region two parts may be more likely to be near each other by accident, rather than because they are corresponding parts of similar objects; the opposite is true in a low density region.

Instead of measuring the similarity of two parts by a distance value D (e.g., Euclidean distance) in the representation space, we will measure their similarity by the probability of obtaining the observed value of D by chance. Specifically, given a particular part p and a second part q , we measure the probability of obtaining a part within distance $D(p,q)$ of p by chance. The chance hypothesis in this case is defined as random selection, with equal probability, of a part from the set of all parts of all example shapes of all classes in a given database of shape examples. The less likely that a value of D less than or equal to $D(p,q)$ would occur by chance, the more similar the two parts are considered to be. In the following, the term "chance probability" will be used as shorthand for "the probability of an event occurring as the result of the chance hypothesis."

It is helpful to compare this type of inference to four others widely used in computer vision.

(1) In classical *statistical hypothesis testing*, a chance hypothesis is rejected or not rejected; in the current method, the chance probabilities are ranked. This enables thousands of hypotheses to be compared. (2) In inference from *nonaccidentalness* [7,13,18,20,28], the chance probabilities of complex events (e.g., matches between two contour segments [20]) have typically been estimated by combining the chance probabilities of simple events (e.g., matches between contour points). In contrast, the present approach estimates the chance probabilities of complex events directly, from the empirical distribution in the knowledge base. As a result, feature combination models such as independence or Markov models are not required. Also, the numerical inaccuracy of combining many probabilities is avoided, enabling our approach to be used in high-dimensional spaces. The examples in this paper have a dimensionality of 200. (3) In *class-based inference*, class membership information is required. By contrast, the current approach does not require class labels on the database, and thus can be used when class membership information is not available. (4) In *kernel-based* methods, a window size parameter must be specified. By contrast, there are no free parameters in the estimation of the chance probability functions.

The basic element of the proposed method is the *chance probability function* (CPF). There is one such function for each part, p , of a shape to be matched. The value of this function at metric distance value D_0 in the representation space is the probability of getting a part within D_0 of p by chance. It is estimated by

$$CP_p(D_0) = | \{ D(p,p') \leq D_0 : p' \in S \} | / | S | , \quad (1)$$

where S is a large set of parts from example shapes of multiple classes, $| \cdot |$ is the set cardinality function, D is a distance measure in a part representation space, and p' is an arbitrary part in S . The CPF is a cumulative probability distribution; the term chance probability function is used to indicate that it is specifically the distribution with respect to the chance hypothesis. The CPFs are computed for each part of the shape to be matched, using a large part database as the set S . Then, for each part p in one shape to be matched and each part q in the other shape to be matched, the CPF of p is used to map the geometric difference $D(p,q)$ to the probability of that

difference, $CP_p(D(p,q))$. The details are available in [24,27]; in particular, the CPFs are represented as look-up tables with a high degree of compression (e.g., 100-to-1 compression), they are linearly interpolated to change them from discrete to continuous, and the chance probabilities are symmetrized:

$$D_{CP}(p,q) = \max \{ CP_p[D(p,q)], CP_q[D(p,q)] \}. \quad (2)$$

The central experiment presented in this paper is a direct comparison of the metric distance $D(p,q)$ with the probabilistic part difference measure $D_{CP}(p,q)$. All other factors – the part generation method, the part representation method, and the matching algorithm – are held constant.

2.2 Part generation and representation

The part correspondence method is independent of the part generation and representation methods used. This paper uses simple, rapidly computed methods. More sophisticated methods may be substituted (e.g., [19]), but we have not found that necessary.

Shapes are represented by closed or open planar contours. We discuss the closed contour case; the open contour case is nearly identical. Let c denote a contour. A set of key points $K = \{K_1, \dots, K_n\}$ is found. In the current implementation, key points are significant positive local maxima and negative local minima of curvature [24] (Fig. 1a). The set of parts S_I is defined to be the set of all intervals between distinct key points on c . Let $[K_i, K_j]$ denote the restriction of c to the oriented interval between key points K_i and K_j . (Note that on a closed contour, $[K_i, K_j] \cup [K_j, K_i] = c$.) There are $n(n-1)$ parts. Most shapes we have examined (about 1,500) have between 10 to 40 key points, so in practical terms $|S_I|$ is effectively bounded.

Each part in S_I is represented in an invariant reference frame by translating, rotating, and scaling the part so that $K_i \rightarrow (0,0)$ and $K_j \rightarrow (1,0)$ (Fig. 1b). Within the invariant reference



Figure 1. (a) Key points are positive local maxima and negative local minima of curvature. (b) An example of a normalized part (from front of dorsal fin to center notch of tail).

frame, every part p is resampled at uniform arc length intervals by α points $\mathbf{x}_1, \dots, \mathbf{x}_\alpha$, where $\mathbf{x}_i = (x_i, y_i)$. A sum of squared distances (SSD) difference measure between parts p, p' is defined by

$$D_{SSD}(p, p') = \sum_{i=1}^{\alpha} (\mathbf{x}_i - \mathbf{x}'_i)(\mathbf{x}_i - \mathbf{x}'_i)^T. \quad (3)$$

Each part can also be viewed as a concatenated vector $\mathbf{v} = (\mathbf{x}_1, \dots, \mathbf{x}_\alpha)$ in a 2α -dimensional part space; then $D_{SSD}(p, p')$ is squared Euclidean distance in that part space. It is important to note that the match is component-by-component vector matching; there is no attempt to solve the point correspondence problem in its usual sense of determining which points in p' are closest to, or correspond best to, points in p [3,5].

2.3 Part correspondence

This section presents the definition of a part correspondence between two shapes and a method for finding that correspondence.

Let S_E denote the set of *elementary parts* of a shape, defined by adjacent key points, and let S_C denote the closure of S_E under set union. Then $S_E \subseteq S_I \subseteq S_C$.

Elementary parts are not generally useful for part correspondence because they are often too simple to match unambiguously; furthermore, corresponding elementary parts might not exist due to contour segmentation differences. Instead, given an elementary part π , we find the part p in the set of all parts containing π , such that p has the best match to a part p' in the other shape. The part p can be thought of as the adaptively selected context for π that can be most confidently matched to a part in the other shape, in the sense of having the least probability of occurring by chance. The part correspondence will be constructed from a set of adaptively selected context parts that cover the set of elementary parts. Since multiple elementary parts can share the same containing parts, redundant part comparisons can be avoided by using the following greedy algorithm: (1) select the best matching pair of parts; (2) delete subparts of that pair; (3) output the portions of the pair that cover unmatched sections of the contour; (4) repeat until no more matches can be performed. We now define the concept and algorithm formally; readers not interested in the details may proceed directly to Section 3.

Let s, s' denote two shapes. Let $\pi \in S_E$ be an elementary part in shape s . Let $S_I(\pi) \subseteq S_I$ denote the set of parts in S_I that contain π . Let δ denote a general part difference measure. Given an elementary part π in the first shape, let $(p(\pi), p'(\pi)) \in S_I(\pi) \times S'_I$ denote the pair of parts that has the minimum difference among all part pairs, such that the first part of the pair contains π :

$$(p(\pi), p'(\pi)) = \underset{(q, q') \in S_I(\pi) \times S'_I}{\operatorname{argmin}} \{ \delta(q, q') \}. \quad (4)$$

The experiments will compare $\delta = D_{SSD}$ (squared Euclidean distance, Eqn. 3) and $\delta = D_{CP}$ (chance probability, Eqn. 2). Let $d(\pi) = \delta(p(\pi), p'(\pi))$ denote the minimizing value in (4).

The complete part correspondence, $C(s, s')$, between shapes s and s' is constructed as follows. Let $k = |S_E|$ denote the number of elementary parts of s . Index the elementary parts in order of increasing $d(\pi)$, breaking ties arbitrarily: i.e., $p(\pi_1), \dots, p(\pi_k)$, with $d(\pi_i) \leq d(\pi_j)$ for $i < j$. Let $C^*(s, s') = \{(p(\pi_i), p'(\pi_i))\}_{i=1}^k$. Generally, $|C^*(s, s')| < k$ since multiple elementary parts π_i may share the same containing part (i.e., $p(\pi_i) = p(\pi_j)$ for some $i \neq j$). Also, note that the parts in $C^*(s, s')$ can overlap on the contour. A set of correspondences $C^{**}(s, s') = \{(r_1, r'_1), \dots, (r_k, r'_k)\}$ that is disjoint in each contour except for shared endpoints can be generated by

$$r_1 = p_1 \text{ and } r_i = p_i - \bigcup_{j=1}^{i-1} p_j \text{ for } 1 < i \leq k, \quad (5)$$

where $p_i = p(\pi_i)$, "-" denotes set difference, and $r_i \in S_C$. The parts r'_i are generated from p'_i in the same way. Redundant parts in $C^*(s, s')$ become empty parts in $C^{**}(s, s')$. By eliminating the redundant parts, we obtain the final *part correspondence* between s and s' , denoted by $C(s, s') = \{(r_1, r'_1), \dots, (r_z, r'_z)\}$, where $z \leq k$. It is important to note that this definition allows parts in one shape to have no match in the other shape. This can occur, for example, when one shape has extra structure not present in the other shape. The algorithm used to generate the part correspondence $C(s, s')$ is given in Fig. 2. Fig. 3 illustrates the algorithm on a simple example.

-
- Notation: i = iteration number.
 Δ = matrix of part-to-part differences between S_i, S'_i .
 C_i = correspondence after iteration i .
 U_i, U'_i = elementary parts for which correspondences have been found so far.
 $E(r)$ = set of elementary parts contained in part r .
1. $i = 0$. $U_i = U'_i = C_i = \emptyset$.
 2. While $U_i \neq S_E$ and $U'_i \neq S'_E$ and Δ is not empty, do:
 - a. $i = i + 1$.
 - b. Find the minimum entry of Δ . Let $p_i \in S_i$ and $p'_i \in S'_i$ be the parts associated with that entry of Δ .
 - c. Compute (r_i, r'_i) from (p_i, p'_i) using (5).
 - d. $C_i = C_{i-1} \cup \{(r_i, r'_i)\}$.
 - e. $U_i = U_{i-1} \cup E(r_i)$. $U'_i = U'_{i-1} \cup E(r'_i)$.
 - f. Remove from further consideration the rows of Δ corresponding to p_i and the parts it contains, and the columns of Δ corresponding to p'_i and the parts it contains.
 3. Output the part correspondences in C_i and the unmatched parts (if any) in s, s' .
-

Figure 2. Algorithm for computing the part correspondence $C(s, s')$ of two shapes s and s' .

3. Experiments

The central thesis of this work is that reliable part correspondences between two shapes can be found directly when a knowledge base is available. The knowledge used is statistical information about the distribution of parts in part space as captured in the chance probability functions. To demonstrate the approach, we will compare it to a classical approach using Euclidean distance in part space.

The database of parts used in the experiments was generated by choosing 10 classes from a




Figure 3. Example of matching process. (a) Input shapes. (b) Key points. (c)-(d) First iteration: best matching part pair (p_1, p'_1) and output pair (r_1, r'_1) , respectively. (e) Second iteration: best matching part pair (p_2, p'_2) after all parts subsumed in first match are removed from consideration. (f) Output pair (r_2, r'_2) , consisting of the portions of the parts in (e) not overlapping the parts in (r_1, r'_1) . (g) Third iteration: best matching part pair (p_3, p'_3) after all parts subsumed in prior matches are removed from consideration. (h) Output pair (r_3, r'_3) , consisting of portions of parts in (g) not overlapping parts (r_1, r'_1) and (r_2, r'_2) .



Figure 4. Two examples from each of the ten object classes used to generate the parts database.

large database of silhouette shapes [16]. Five shapes were randomly selected from each class. (Fig. 4 shows two examples of each class.) The mirror reflections of the shapes were also used, to yield a total of 100 shapes. The part generation procedure described in Section 2.2 was applied to each shape, yielding a database of 19,220 parts. At runtime, the part database was used to compute the chance probability function CP_p for each part p in the two input shapes to be matched. In all of the experiments, the input shapes were different from those used to generate the part database.

The part correspondences were generated using the algorithm in Fig. 2 with two shape difference measures: (1) squared Euclidean distance between the part vectors in 2α -dimensional space ($\delta = D_{SSD}$), and (2) chance probabilities ($\delta = D_{CP}$). To handle reflections, both $C(s, s')$ and $C(R(s), s')$ were computed, where $R(s)$ is the reflection of s ; the version with the lower total match score was used.

Fig. 5a shows two new elephant shapes (i.e., not among those used to generate the part database). Fig. 5b shows the key points that define the endpoints of extracted parts. When $\delta = D_{SSD}$, the part correspondence consists of 20 part pairs. The best four matches are shown in Fig. 5c. All four of these matches are incorrect. When $\delta = D_{CP}$, the complete part correspondence consists of four part pairs, shown in Fig. 5d. Each of the part correspondences is correct. With




Figure 5. (a) Two previously unseen elephant shapes. (b) Detected key points. (c) The four best part matches using D_{SSD} . Each is incorrect. (d) The complete part correspondence found using D_{CP} . All correspondences are correct.




Figure 6. (a) Two previously unseen ray shapes. (b) Detected key points. (c) The four best part matches using D_{SSD} . The first three are incorrect. (d) The complete part correspondence found using D_{CP} . All correspondences are correct. Notice that in (d) the extra structure in the right ray's tail is correctly found not to match any part of the left ray.

squared Euclidean distance, the matched parts are smaller, less distinctive, and often incorrectly corresponded. With chance probabilities, the parts are larger, more complex, and correctly corresponded, even though some of the corresponding parts are quite different geometrically.

It must be emphasized that the method does not use class information in any way. All of the information resides in the CPFs which are computed from a single set of parts from all classes. The parts in the database are not labeled by class nor by the object they came from.

In the introduction we claimed that part correspondence could be used to facilitate point correspondence by subdividing a difficult point matching problem into easier subproblems. To illustrate, consider the elephant pair in Fig. 5. The thin plate spline - robust point correspondence (TPS-RPM) algorithm of [5] did not converge on this pair because the shapes are complex and geometrically dissimilar. (It was tested with a number of parameter values.) We then applied the TPS-RPM algorithm separately to each pair of parts found by the part correspondence algorithm and concatenated the solutions. The result, shown in Fig. 8, is a high-quality point correspondence over most of the shape.

Fig. 6a shows two previously unseen ray shapes. Fig. 6b shows the key points. When $\delta = D_{SSD}$, the part correspondence consists of 10 part pairs. The best four matches are shown in Fig. 6c. Three of these are incorrect. (The shapes are matched under reflection because the rays' tails curve in opposite directions, so corresponding parts should occur on opposite sides.) When $\delta = D_{CP}$, the complete correspondence consists of four part pairs (Fig. 6d). Each of them is correct, including the null match for the structure on the right ray's tail that does not have a counterpart




Figure 7. Part correspondence in the presence of repetitive structure. (a) Two previously unseen tree shapes. (b) Key points. (c-e) The complete part correspondence found using D_{CP} .

on the left ray's tail. This example also demonstrates the rotation invariance of the method.

Fig. 7 shows similar results for two previously unseen tree shapes. This example shows correct matching of parts even in the presence of repetitive structure. The method implicitly adapts the context to obtain unambiguous matches since these are less likely to occur by chance.

Novel shape pairs of the remaining seven object classes used to generate the parts database produced similar results. Although the chance probability method occasionally produced incorrect part correspondences, and SSD occasionally produced correct correspondences, the chance probability model always did as well as or better (usually much better) than SSD.

The time required to compute a complete part correspondence between two shapes varied according to the complexity of the shapes. Times ranged from 6 to 51 s with an average of 22 s on a 1.8 Ghz Pentium PC running interpreted Matlab. Much of the computational time is spent computing the CPFs, currently done at runtime. Speeding up the method by precomputing information about the parts database is one area for future research.

We also tested the possibility that a modified SSD measure normalized by the arc lengths of the parts would perform better than D_{SSD} . A second experiment (not shown) demonstrated that normalized SSD also produced many incorrect part correspondences.



Figure 8. The output of the Chui-Rangarajan TPS-RPM point matcher [5] applied part-wise to the part pairs found by the part correspondence method of this paper. For display purposes, the dense point correspondences are subsampled in the figure.




Figure 9. (a) Two shapes from a novel class. (b) Key points. (c-e) The complete part correspondence found using D_{CP} . The part correspondences are correct even though no hand shapes were used to generate the parts database.

Finally, Figure 9 shows the result for two shapes (hands) that were not from any of the object classes used to generate the part database. The part correspondence found is correct, presumably because the part database provides enough information about the general distribution of parts in part space to enable the matches between the parts of the hands to be disambiguated. As before, SSD results in an incorrect part correspondence.

4. Discussion

Many knowledge-based approaches use either models of classes or examples labeled by class (e.g., [6,22]). The present approach is different in that no information about classes is used: the knowledge consists of the density of parts in part space across multiple classes. That this relatively low-level knowledge can be so effective in disambiguating matches is the most important demonstration of this paper.

We compared a geometric difference measure D_{SSD} and a probabilistic difference measure D_{CP} using the same algorithm. The results, together with earlier results on shape retrieval and classification using CPFs [24,27], suggest that context in the representation space (provided by D_{CP}) may be more important than pure geometric matching (D_{SSD}) in determining shape similarity.

For a given elementary part, the method adaptively chooses the size of the context in the shape that maximizes the nonaccidentalness, and hence the confidence, of the match. This represents an advance over matching methods that disambiguate matches using global context [2] or local contexts defined by a fixed number of feature points [12]. The adaptive behavior appears automatically when chance probabilities are substituted for geometric differences in the

greedy algorithm. The chance probabilities provide the connecting mechanism whereby the context of a part in the representation space determines the context of that part within its shape.

It is important to note that the present method does not solve a point correspondence problem between parts. The matching of parts to find D_{SSD} (which is then used to find D_{CP} using the CPF lookup tables) is fixed, component-by-component vector matching: there is no attempt to solve the correspondence problem in the usual sense of determining which point of one part best matches which point of another part. (I.e., the 'correspondence' is independent of the data.)

The present work should also not be confused with methods that extend the point correspondence concept from one-to-one matches to many-to-one or many-to-many matches [5]. We suggest, similarly to [4], that it is not necessarily meaningful to talk about point correspondences in some cases. Consider again the two front legs of the elephants in the rightmost pair in Fig. 5d. These parts are geometrically so different that, except at certain distinguished points, any proposed point correspondence amounts to choosing one interpolation scheme over another. However, although the geometries of these parts correspond loosely at best, the parts themselves correspond unambiguously.

It is also important to note that the method does not explicitly model or match part relationships between the two shapes. An intriguing area of future research is to explore the use of part correspondences to reduce the size of the search space in structural matching approaches.

Another area of future research is the development of part-based class models, by extending the method from two-shape to multi-shape part correspondences, analogously to the use of multi-shape point correspondences to build point-based class models [14].

5. Conclusion

This paper advocated two ideas: (1) part correspondence as a new type of shape comparison, and (2) knowledge-based shape matching. The paper presented a class-free model in which the knowledge used to match the parts of two shapes is embodied in chance probability functions computed from a large set of parts of many different shapes. The paper also presented a method

for finding a set of part correspondences between two shapes. A side-by-side comparison of Euclidean distance with the probabilistic difference measure showed that the power of the method is due to the chance probability model. Part correspondence is a potentially valuable form of shape matching, both in its own right and as a complement to geometric and structural matching.

Acknowledgements

The author thanks Drs. Latecki [16] and Mokhtarian [19] for making contour data files available, and Drs. Chui and Rangarajan [5] for the TPS-RPM program.

References

- [1] C. Aslan and S. Tari, An Axis Based Representation for Recognition, *International Conf. on Computer Vision (ICCV)*, 1339-1346, 2005.
- [2] S. Belongie, J. Malik, and J. Puzicha, Shape Matching and Object Recognition Using Shape Contexts, *IEEE Trans. Pattern Analysis and Machine Intelligence* **24**(4), 509-522, 2002.
- [3] P. Besl and N. D. McKay, A Method for Registration of 3-D Shapes, *IEEE Trans. Pattern Analysis and Machine Intelligence* **14**(2), 239-256, 1992.
- [4] S. Carlsson, Order structure, correspondence and shape based categories, *Int'l Workshop on Shape, Contour and Grouping*, Springer Lecture notes in Computer Vision, 58-71, 1999.
- [5] H. Chui and A. Rangarajan, A New Point Matching Algorithm for Non-Rigid Registration, *Computer Vision and Image Understanding* **89**, 114-141, 2003.
- [6] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham, Active Shape Models - Their Training and Application, *Computer Vision and Image Understanding* **61**(1) 38-59, 1995.
- [7] A. Desolneux, L. Moisan, and J.-M. Morel, Meaningful Alignments, *International Journal of Computer Vision* **40**(1), 7-23, 2000.
- [8] P. F. Felzenszwalb and D. P. Huttenlocher, Efficient Matching of Pictorial Structures, *IEEE Conf. Computer Vision and Pattern Recognition*, II:66-73, 2000.
- [9] Y. Gdalyahu and D. Weinshall, Flexible Syntactic Matching of Curves and Its Application to Automatic Hierarchical Classification of Silhouettes, *IEEE Trans. Pattern Analysis and Machine Intelligence* **21**(12), 1312-1328, 1999.
- [10] D. Geiger, T. Liu, and R. V. Kohn, Representation and Self-Similarity of Shapes, *IEEE Trans. Pattern Analysis and Machine Intelligence* **25**(1), 86-99, 2003.
- [11] A. Ghosh and N. Petkov, Robustness of Shape Descriptors to Incomplete Contour Representations, *IEEE Trans. Pattern Analysis and Machine Intelligence* **27**(11), 1793-

- 1804, 2005.
- [12] C. Grigorescu and N. Petkov, Distance Sets for Shape Filters and Shape Recognition, *IEEE Trans. Image Processing* **12**(10) 1274-1286, 2003.
- [13] W. E. L. Grimson and D. P. Huttenlocher, On the Verification of Hypothesized Matches in Model-Based Recognition, *IEEE Trans. Pattern Analysis and Machine Intelligence* **13**(12) 1201-1213, 1991.
- [14] A. Hill, C. J. Taylor, and A. D. Brett, A Framework for Automatic Landmark Identification Using a New Method of Nonrigid Correspondence, *IEEE Trans. Pattern Analysis and Machine Intelligence* **22**(3) 241-251, 2000.
- [15] L. Latecki and R. Lakämper, Shape Similarity Measure Based on Correspondence of Visual Parts, *IEEE Trans. Pattern Analysis and Machine Intelligence* **22**(10) 1185-1190, 2000.
- [16] L. Latecki, R. Lakämper, and U. Eckhardt, Shape Descriptors for Non-rigid Shapes with a Single Closed Contour. *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 424-429, 2000.
- [17] L. Latecki, R. Lakaemper, and D. Wolter, Optimal Partial Shape Similarity, *Image and Vision Computing* **23**(2) 227-236, 2005.
- [18] D. G. Lowe, Three-Dimensional Object Recognition from Single Two-Dimensional Images, *Artificial Intelligence* **31**(3) 355-395, 1987.
- [19] F. Mokhtarian, S. Abbasi, and J. Kittler, Efficient and Robust Retrieval by Shape Content through Curvature Scale Space, in A. Smeulders and R. Jain (Eds.), *Image Databases and Multi-Media Search*, World Scientific, New Jersey, pp. 51-58, 1997.
- [20] C. F. Olson and D. P. Huttenlocher, Automatic Target Recognition by Matching Oriented Edge Pixels, *IEEE Trans. Image Processing* **6**(1) 103-113, 1997.
- [21] T. B. Sebastian, P.N. Klein, and B. B. Kimia, Recognition of Shapes by Editing Their Shock Graphs, *IEEE Trans. Pattern Analysis and Machine Intelligence* **26**(5) 550-571, 2004.
- [22] K. B. Sun and B. J. Super, Classification of Contour Shapes Using Class Segment Sets, *IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, San Diego, CA, 727-733, 2005.
- [23] B. J. Super, Fast Retrieval of Isolated Visual Shapes, *Computer Vision and Image Understanding* **85**(1) 1-21, 2002.
- [24] B. J. Super, Learning Chance Probability Functions for Shape Retrieval or Classification, *Proc. IEEE Workshop on Learning in Computer Vision and Pattern Recognition*, Washington DC., June 2004.
- [25] B. J. Super, Knowledge-Based Part Correspondence, University of Illinois at Chicago

Computer Vision and Robotics Laboratory Technical Report CVRL TR 2004-07-001, July 2004.

- [26] B. J. Super, Shape Part Correspondence Using Chance Probability Functions, *Proc. of the 2006 International Conference on Image Processing, Computer Vision, and Pattern Recognition*, Las Vegas, NV, 269-275, 2006.
- [27] B. J. Super, An Effective and Fast One-to-Many Shape Matcher, *International Journal of Pattern Recognition and Artificial Intelligence*, in press.
- [28] A. Witkin and J. Tenenbaum, On the Role of Structure in Vision, in J. Beck, B. Hope, and A. Rosenfeld, (Eds.), *Human and Machine Vision*, NY: Academic Press, pp. 481-543, 1983.
- [29] S. C. Zhu and A. L. Yuille, FORMS: A Flexible Object Recognition and Modeling System, *International Journal of Computer Vision* **20**(3) 187-212, 1996.