

On Aligning Curves

Thomas B. Sebastian, Philip N. Klein, and
Benjamin B. Kimia

Abstract—We present a novel approach to finding a correspondence (alignment) between two curves. The correspondence is based on a notion of an alignment curve which treats both curves symmetrically. We then define a similarity metric based on the alignment curve using two intrinsic properties of the curve, namely, length and curvature. The optimal correspondence is found by an efficient dynamic-programming method both for aligning pairs of curve segments and pairs of closed curves, and is effective in the presence of a variety of transformations of the curve. Finally, the correspondence is shown in application to handwritten character recognition, prototype formation, and object recognition, and is potentially useful in other applications such as registration and tracking.

Index Terms—Curve alignment, recognition, dynamic programming, prototypes, correspondence.

1 INTRODUCTION

THIS paper presents a novel method to find the optimal correspondence or *alignment* between 2D curves using their intrinsic properties. This alignment has served as a key element in many applications such as object recognition based on silhouettes [1], [2], [3], handwritten character recognition [4], [5], [6], tracking [7], [8], etc. While existing curve alignment/matching approaches typically take advantage of constraints pertaining to their particular domain of application, we approach this problem generically, keeping in mind that domain-specific constraints can be added subsequently. First, we review current approaches for curve alignment, discuss how we extend current techniques, and present an overview of our approach.

Current curve alignment methods can be classified into two categories: methods based on rigid transformations [1], [9], [10], and those based on nonrigid deformations [7], [8], [11], [12]. Methods based on *rigid transformations* rely on matching feature points by finding the optimal rotation, translation, and scaling parameters [1], [10]. Since these methods assume that the outlines can be aligned by a rigid transformation, they are sensitive to articulations, deformations of parts, occlusion, and other variations in the object form. Methods based on nonrigid *deformations* model articulation and other deformations by finding the mapping from one curve to another that minimizes a performance functional consisting of “stretching” and “bending” energies [7], [8], [11], [12]. The minimization problem in the discrete domain is transformed into one of matching shape signatures with curvature, bending angle, or absolute orientation as attributes [5], [2], [3]. These methods suffer from one or more of the following drawbacks: asymmetric treatment of the two curves, sensitivity to sampling of the curves, lack of rotation and scaling invariance, and sensitivity to occlusion and articulations. We explore an extension of these approaches that addresses some of these issues.

Our approach to finding the optimal correspondence between 2D curves relies on using the *intrinsic* properties of the curves in an energy minimization framework as in [7], [12]. However, a key

problem in this framework is the *asymmetric treatment* of the curves being matched. Tagare [8] handles this issue by introducing a bimorphism and searching in the space of *pairs* of functions. Observing a redundancy in this search, we introduce the notion of an *alignment curve* which ensures a symmetric treatment by searching in the space of a single function, Section 2. This naturally leads to an “edit distance” metric. Efficient algorithms for finding the optimal alignment using dynamic programming are described in Section 3 for matching both curve segments and closed curves. Section 4 illustrates the generic nature of the proposed curve alignment framework in several applications including handwritten character recognition, prototype formation, and object recognition.

2 CURVE ALIGNMENT

This section discusses the mathematical formulation of the problem of aligning two curves. We first consider the case of aligning two curve segments and then extend this to align two closed curves.

2.1 Deformation-Based Approach

Cohen et al. [7] pioneered the deformation-based approach to curve matching. The basic premise of their approach was to match high curvature points along the curves, while maintaining a smooth displacement field. The problem was cast in terms of minimizing an energy functional penalizing “bending” and “stretching” in a physical analogy similar to the one used in formulating active contours [13]. Younes [12] formalized the approach. Tagare [8] pointed out the inherent asymmetry in the treatment of the two curves and proposed a “bimorphism” that ensures a symmetric formulation. We now give the specifics of the problem formulation.

Denote the curve segments to be matched by $C(s) = (x(s), y(s))$, $s \in [0, L]$, and $\bar{C}(\bar{s}) = (\bar{x}(\bar{s}), \bar{y}(\bar{s}))$, $\bar{s} \in [0, [\bar{L}]]$, where s is arc length, x and y are coordinates of each point, L is length of C , and each is similarly defined for \bar{C} . A central premise of this approach is that the “goodness” of the optimal match is the sum of the “goodness” of the matches between corresponding subsegments. This allows an energy functional to convey the goodness of a match as a function of the correspondence or alignment of the two curves as proposed earlier [7], [12]. Let a mapping $g: [0, L] \rightarrow [0, [\bar{L}]$, $g(s) = \bar{s}$, represent an alignment of the two curves. Specifically, Cohen et al. [7] compare the displacement velocities and bending energies in the form of

$$\mu[g] = \int_C \left| \frac{\partial}{\partial s} (\bar{C}(\bar{s}) - C(s)) \right|^2 ds + R \int_C (\kappa(s) - \bar{\kappa}(\bar{s}))^2 ds, \quad (1)$$

where κ and $\bar{\kappa}$ are the curvatures along the curves C and \bar{C} , respectively, R is a parameter, and $\bar{s} = g(s)$. Younes [12] uses a similar functional. Note that these methods are not invariant to the relative *rotation* of the two curves. Hence, the optimal relative orientation must also be found. Also, in Cohen’s method, the cost of the optimal deformation is not invariant to the way the curves are sampled, due to the curvature comparison term. We now formulate the problem in an intrinsic manner which addresses both issues.

2.2 Intrinsic Formulation

Definition 1. Let $C|_{[s_1, s_2]}$ denote the portion of the curve from s_1 to s_2 , and let $g|_{([s_1, s_2], [\bar{s}_1, \bar{s}_2])}$ denote the restriction of the mapping g to $[s_1, s_2]$, where $\bar{s}_i = g(s_i)$, $i = 1, 2$. Define a measure μ on this alignment function,

$$\mu[g] |_{([s_1, s_2], [\bar{s}_1, \bar{s}_2])} : g |_{([s_1, s_2], [\bar{s}_1, \bar{s}_2])} \rightarrow \mathfrak{R}^+,$$

to denote the cost of deforming $C|_{[s_1, s_2]}$ to $\bar{C}|_{[\bar{s}_1, \bar{s}_2]}$.

- T.B. Sebastian and B.B. Kimia are with the Division of Engineering, Brown University, Providence, RI 02912. E-mail: {tbs, kimia}@lems.brown.edu.
- P.N. Klein is with the Department of Computer Science, Brown University, Providence, RI 02912. E-mail: klein@cs.brown.edu.

Manuscript received 16 Feb. 2001; revised 27 Dec. 2001; accepted 21 Mar. 2002.

Recommended for acceptance by D. Jacobs.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number 113646.




Fig. 1. (a) The cost of deforming an infinitesimal segment AB to segment \overline{AB} , when the start points and their tangents are aligned ($A = \overline{A}$, $\overline{T}_A = \overline{T}_{\overline{A}}$), is related to the distance $B\overline{B}$, and is defined by $|d\overline{s} - ds| + R|d\overline{\theta} - d\theta|$. (b) The formulation using an alignment curve allows for a finite segment from one curve to be aligned with a single point on one curve, thus allowing for the curve segment deletion or addition.

We restrict this measure μ to one which satisfies an *additivity property*

$$\mu[g] |_{([s_1, s_3], [\overline{s}_1, \overline{s}_3])} = \mu[g] |_{([s_1, s_2], [\overline{s}_1, \overline{s}_2])} + \mu[g] |_{([s_2, s_3], [\overline{s}_2, \overline{s}_3])}, \quad (2)$$

$$\forall s_1 \leq s_2 \leq s_3 \in [0, L], \forall \overline{s}_1 \leq \overline{s}_2 \leq \overline{s}_3 \in [0, \overline{L}],$$

where $\overline{s}_i = g(s_i)$, $i = 1, 2, 3$. This property implies that the match process can be decomposed into a number of smaller matches, and can be written as a functional

$$\mu[g] |_{([0, L], [0, \overline{L}])} = \int_0^L \mu[g] |_{([s, s+ds], [g(s), g(s+ds)])} ds. \quad (3)$$

The optimal match is then given by

$$g^* = \arg \min_g \mu[g] |_{([0, L], [0, \overline{L}])}.$$

Consider two infinitesimal curve segments $C|_{[A, B]}$ and $\overline{C}|_{[\overline{A}, \overline{B}]}$ of lengths ds , $d\overline{s}$, and curvatures κ , $\overline{\kappa}$, respectively. Since we compare the intrinsic aspects of these curve segments, we align their start points A and \overline{A} and their respective tangents \overline{T}_A and $\overline{T}_{\overline{A}}$. Fig. 1a. The cost of matching the infinitesimal curve segments is the degree by which the endpoints B and \overline{B} differ, which can be formulated as

$$\mu[g] |_{([s_1, s_2], [ds, [\overline{s}_1, \overline{s}_2] + d\overline{s}])} = |d\overline{s} - ds| + R|d\overline{\theta} - d\theta|, \quad (4)$$

where R is a constant.¹ Then, the resulting functional is given by

$$\mu[g] = \int_0^L \left[\left| \frac{d\overline{s}}{ds} - 1 \right| + R \left| \frac{d\overline{\theta}(\overline{s})}{d\overline{s}} \frac{d\overline{s}}{ds} - \frac{d\theta(s)}{ds} \right| \right] ds \quad (5)$$

$$= \int_0^L [|g'(s) - 1| + R|\overline{\kappa}(g(s))g'(s) - \kappa(s)|] ds,$$

where the first term penalizes “stretching” and the second term penalizes “bending.” This is the same as the linear cost model of Basri et al. [11].² Observe that, while this formulation is symmetric in form [11], the explicit dependence on the alignment function g

1. R is a scale constant depending on the average sample size along the curve. For all examples in this paper, we set $R = 10$. We have seen experimentally that the alignment is relatively insensitive to the choice of R .

2. We refer the reader to Basri et al. [11] for a discussion of different nonlinear norms which were motivated by the extent certain desirable constraints were met. These norms can also be used in the curve alignment framework.

makes it inherently asymmetric [8]. Specifically, algorithms which are based on differentiable mapping of one curve to the other are asymmetric, as the inverse map may not be one-to-one or differentiable. Tagare [8] instead proposes a “bimorphism,” which diffeomorphically maps the pair of curves to be matched. Specifically, he formulates a cost function that minimizes differences in local length and orientation changes $|d\overline{s} - ds|$ and $|d\overline{\theta} - d\theta|$ along each differential segment of this curve, and seeks a *pair* of functions ϕ_1 and ϕ_2 (elements of the bimorphism) which optimize this cost functional.

2.3 Symmetric Formulation: Alignment Curve

We approach this asymmetry issue by observing that the formulation based on (5) does not allow the mapping of a single point in the first curve to a segment in the second curve. This is because the notion of an alignment is captured by a (univalued) function g . To alleviate this difficulty, we revise the formulation. We reconsider an alignment between two curves as a pairing of two particles, one on each curve traversing their respective paths monotonically, but with finite stops allowed. This alignment can be specified in terms of two functions h and \overline{h} relating arc length along C and \overline{C} , s and \overline{s} , respectively, to a newly defined curve parameter ξ , i.e., $s = h(\xi)$ and $\overline{s} = \overline{h}(\xi)$. When h is invertible, we have $\overline{s} = \overline{h}(h^{-1}(s)) = \overline{h} \circ h^{-1}(s)$, which allows for the use of an alignment function, $g = \overline{h} \circ h^{-1}$, as before. However, when h is not invertible, i.e., when the first particle stops along the first curve for some finite time, g is not defined.

While this formulation allows for a symmetric treatment of the curves, note that a superfluous degree of freedom is introduced, as



Fig. 2. Some characters differ in curvature only at a few points and have small edit distance. In this figure, “6” and “U” differ in curvature only at the few points that are highlighted.

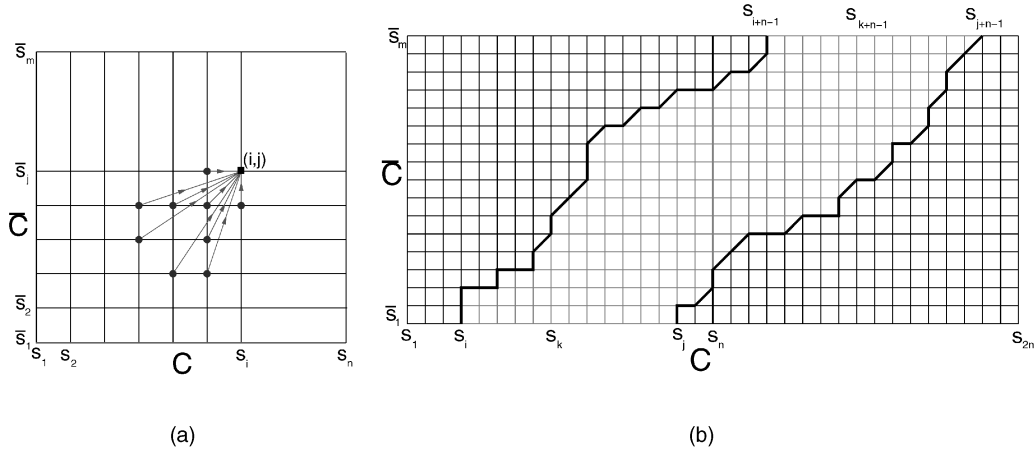


Fig. 3. (a) This figure illustrates the template that is used to find the edit distance of curve segments using dynamic programming. Discrete samples along the curves are the axes. The entry at (i, j) represents $d(i, j)$. To update the cost at (i, j) (black square), we limit the choices of the k and l in (11) so that only costs at a limited set of points (gray dots) are considered. (b) This figure illustrates the grid used by the dynamic-programming method to compute the optimal alignment curve for closed curves. The first curve C is repeated. If the bold curves are optimal alignment curves from s_i and s_j , then the alignment curve from s_k for $i < k < j$ does not cross the bold curves, so the search can be restricted to the light gray area in the middle (between the bold curves).

in [8], because different traversals h and \bar{h} may give rise to the same alignment. While Tagare [8] treats this degree of redundancy by simply searching a larger space of two functions, we remove this additional degree of redundancy by introducing the notion of an *alignment curve*, α , with coordinates h and \bar{h}

$$\begin{aligned} \alpha(\xi) &\triangleq (h(\xi), \bar{h}(\xi)), \quad \xi \in [0, \tilde{L}], \\ \alpha(0) &= (0, 0), \quad \alpha(\tilde{L}) = (L, \bar{L}), \end{aligned} \quad (6)$$

where ξ is arc-length along the alignment curve and \tilde{L} is its length. The alignment curve can now be specified by a *single function*, namely, $\psi(\xi)$, $\xi \in [0, \tilde{L}]$, the angle between the tangent to the curve and the x -axis. The coordinates h and \bar{h} can then be obtained by integration

$$\begin{cases} h(\xi) = \int_0^\xi \cos(\psi(\eta)) d\eta, \quad \xi \in [0, \tilde{L}]. \\ \bar{h}(\xi) = \int_0^\xi \sin(\psi(\eta)) d\eta, \end{cases} \quad (7)$$

Note that ψ is constrained by *monotonicity* ($h' \geq 0$ and $\bar{h}' \geq 0$) to lie in $[0, \frac{\pi}{2}]$. The alignment between C and \bar{C} is then fully represented by a single function ψ . The goodness of the match between C and \bar{C} can now be rewritten in terms of ψ . First, if $h' \neq 0$ and $\bar{h}' \neq 0$ for $\xi \in [\xi_1, \xi_2]$, then $g = \bar{h} \circ h^{-1}$ is well defined. Using (3) and (4), we get

$$\begin{aligned} \mu[\psi] |_{[\xi_1, \xi_2]} &= \int_{\xi_1}^{\xi_2} \left[\left| \frac{dh}{d\xi} - \frac{d\bar{h}}{d\xi} \right| + R \left| \frac{d\theta ds}{ds d\xi} - \frac{d\bar{\theta} d\bar{s}}{d\bar{s} d\xi} \right| \right] d\xi \\ &= \int_{\xi_1}^{\xi_2} \left[|\cos(\psi) - \sin(\psi)| \right. \\ &\quad \left. + R |\kappa(h) \cos(\psi) - \bar{\kappa}(\bar{h}) \sin(\psi)| \right] d\xi. \end{aligned} \quad (8)$$

Second, consider that one of h' or \bar{h}' is zero at a point, say $h'(\xi) = 0$, implying that $h(\xi)$ maps to a corresponding interval $[\bar{h}(\xi), \bar{h}(\xi + d\xi)]$. This cost of mapping a point to an interval is defined by enforcing continuity of the cost with deformations: consider the cost of mapping the interval $[h(\xi), h(\xi + d\xi)]$ to the interval $[\bar{h}(\xi), \bar{h}(\xi + d\xi)]$ as $dh = h(\xi + d\xi) - h(\xi) \rightarrow 0$, i.e., as $\psi \rightarrow \frac{\pi}{2}$, or as $\cos(\psi) \rightarrow 0$ in (8). Similarly, the case where an interval in the first curve is mapped to a point in the second curve, is the limiting case of $\psi \rightarrow 0$, or $\sin(\psi) \rightarrow 0$ in (8). Since

both limits are well-defined, by using (8) the overall cost of the alignment $\mu[\psi]$ is well defined in all cases, and the optimal alignment is computed as

$$\begin{cases} \psi^* = \operatorname{argmin}_{\psi} \mu[\psi]_{[0, \tilde{L}]} \\ 0 \leq \psi \leq \frac{\pi}{2}, \int_0^{\tilde{L}} \cos(\psi) d\xi = L, \text{ and } \int_0^{\tilde{L}} \sin(\psi) d\xi = \bar{L}. \end{cases} \quad (9)$$

2.4 Distance between Curves

In many computer vision applications including object recognition and handwritten character recognition, finding the similarity between curves, or equivalently a dissimilarity distance between them, is a key element. We now use the alignment computed between two curves to define a distance measure between two curves, which is naturally defined from the cost of deformations:

Definition 2. The edit distance between two curve segments C and \bar{C} is the cost of the optimal alignment of the two curves, $d(C, \bar{C}) = \mu[\psi^*]$.

Since the metric property of a shape distance measure³ is potentially useful in several applications, e.g., in indexing into large databases, symmetry and triangle inequality have been used to efficiently organize the database organization [16], [17], [18], [19], we now show that d is a metric.

Proposition 1. The edit distance is a metric, i.e., d satisfies the following

1. $d(C_1, C_1) = 0$,
2. $d(C_1, C_2) = d(C_2, C_1)$, and
3. $d(C_1, C_3) \leq d(C_1, C_2) + d(C_2, C_3)$.

Proof. 1 and 2 follow directly from the definition of d . To prove 3, let ψ_{ij}^* , $i, j = 1, 2, 3$, specify the optimal alignment between C_i and C_j . Let $[s_1, s_1 + ds_1]$ be an arbitrary interval on C_1 , let $[s_2, s_2 + ds_2]$ be the corresponding intervals on C_2 under alignment ψ_{12}^* , and let $[s_3, s_3 + ds_3]$ be the interval corresponding to $[s_2, s_2 + ds_2]$ on C_3 under alignment ψ_{23}^* . Denote the composed correspondence between C_1 and C_3 which takes $[s_1, s_1 + ds_1]$ to $[s_3, s_3 + ds_3]$ as $\hat{\psi}_{13} = \psi_{12}^* \circ \psi_{23}^*$. Using the inequalities $|df_1 - df_3| \leq |df_1 - df_2| + |df_2 - df_3|$, where $f = s$ or θ , we have

3. There may be instances where human shape comparison does not observe metric properties [14], [15].




Fig. 4. Running time on an SGI Indigo II (195 MHz) for matching curve segments (top) and closed curves (bottom).

$$|ds_1 - ds_3| + R|d\theta_1 - d\theta_3| \leq |ds_1 - ds_2| + R|d\theta_1 - d\theta_2| + |ds_2 - ds_3| + R|d\theta_2 - d\theta_3|,$$

which, when integrated, leads to $\mu[\hat{\psi}_{13}] \leq \mu[\psi_{12}^*] + \mu[\psi_{23}^*]$. Now, using the fact that ψ_{13}^* specifies the optimal alignment of C_1 and C_3 , we have $\mu[\psi_{13}^*] \leq \mu[\hat{\psi}_{13}]$ or $\mu[\psi_{13}^*] \leq \mu[\psi_{12}^*] + \mu[\psi_{23}^*]$, which together with the definition of d completes the proposition. \square

Definition 3. The edit distance between two closed curves is the minimum cost of matching the open curve segments starting at any pair of arbitrary points P and \bar{P} on C and \bar{C} , respectively, and terminating there after having traversed each entire curve, i.e., $d_{\text{closed}}(C, \bar{C}) = \min_{[s_1, \bar{s}_1]} d(C_{[s_1, s_1+L]}, \bar{C}_{[\bar{s}_1, \bar{s}_1+\bar{L}]})$, where L and \bar{L} are the lengths of the closed curves C and \bar{C} , respectively.

The computation of edit distance can be sensitive to the choice of parameters as noted in [2] and also to the manner the curves are sampled. Moreover, computing the normalized edit distance is a nontrivial task [20]. The edit distance, however, has the advantage of working well in the presence of articulation and deformation of parts, because of the use of intrinsic properties. This is critical for object recognition. However, for handwritten character recognition this can also present a drawback, Fig. 2: edit distance penalizes local differences in lengths and curvatures, and some character pairs like “6” and “U,” which differ in curvature only at a few different points, may appear similar. In such cases, global properties must also be taken into account and, thus, the Euclidean distance between corresponding points has been used as a distance measure: The similarity transformation that minimizes the least-squares distance between the corresponding points is found [21], and the average distance between corresponding points is used as the Euclidean distance measure between curves [22], [2]. We will use either the edit distance or the Euclidean distance as is appropriate to the underlying application.




Fig. 5. This figure illustrates the effects of differential sampling. (a) and (b) show the optimal correspondence (dotted lines) for two equal-length line segments (sampled at black dots). Recall that the number of segments that can be merged and then matches is limited to three. (a) When sampling rates are roughly the same, the correspondence and the cost are marginally affected. (b) When the curves are sampled very differently the optimal correspondence is significantly affected. (c) We alleviate the effect of sampling in a postprocessing stage. The points along the optimal alignment curve $\alpha_l, l = k-1, k, k+1, k+2$ are denoted by black squares. In the postprocessing step, these alignment curve points are allowed to “relax” along horizontal and vertical lines. In this case, α_k can become any one of the gray dots.

3 AN ALGORITHM FOR FINDING THE OPTIMAL ALIGNMENT CURVE

This section describes a dynamic-programming algorithm for finding the optimal alignment curve α^* for any two curves C and \bar{C} . Dynamic programming has been used for curve matching [23], [24], [3], [2] deformable-template-based segmentation [25], and in speech recognition [26]. In order to use dynamic programming to find the optimal alignment, the distance function has to satisfy the *optimal substructure* property [27], which we show first.

Proposition 2. The edit distance has optimal substructure, i.e., for any $\xi_1 < \xi_2 < \xi_3$, $s_i = h^*(\xi_i)$, $\bar{s}_i = \bar{h}^*(\xi_i)$, $i = 1, 2, 3$, we have

$$d(C_{[s_1, s_3]}, \bar{C}_{[\bar{s}_1, \bar{s}_3]}) = d(C_{[s_1, s_2]}, \bar{C}_{[\bar{s}_1, \bar{s}_2]}) + d(C_{[s_2, s_3]}, \bar{C}_{[\bar{s}_2, \bar{s}_3]}). \quad (10)$$

The basic idea is that under the optimal alignment, the alignment of aligned subsegments also has to be optimal. If that is not the case, the nonoptimal alignments of the subsegments can be replaced by the optimal one to get an alignment with a lower total cost. The detailed proof is omitted due to space constraints [28]. We now describe how the optimal alignment curve α^* is found.

The alignment curve of C and \bar{C} is a curve in the 2D plane whose axes are specified by the curve segments C and \bar{C} . We discretize C and \bar{C} at samples s_1, s_2, \dots, s_n and $\bar{s}_1, \bar{s}_2, \dots, \bar{s}_m$, respectively, Fig. 3a. The alignment curve is then represented by a sequence of N points $(\alpha_1, \dots, \alpha_N)$, where

$$\begin{cases} \alpha_k = (s_{i_k}, \bar{s}_{j_k}), & i_k \in \{1, \dots, n\}, j_k \in \{1, \dots, m\}, \\ & k = 1, \dots, N, \\ \alpha_1 = (s_1, \bar{s}_1), \\ \alpha_N = (s_n, \bar{s}_m). \end{cases}$$

That we have constrained the match at each respective endpoints of the two curve segments is not a restriction since the match also considers all possible deletions of segments from the beginning and the end of each curve.

Let $d(i, j)$ denote the cost of matching the discrete curve segments $C|_{[s_1, s_i]}$ and $\bar{C}|_{[\bar{s}_1, \bar{s}_j]}$; let $\delta([k, i], [l, j])$ denote the cost of matching subsegments $C|_{[s_k, s_i]}$ and $\bar{C}|_{[\bar{s}_l, \bar{s}_j]}$. The optimal substructure property of the distance function (10) allows us to write

$$d(i, j) = \min_{k, l} [d(i - k, j - l) + \delta([i - k, i], [j - l, j])], \quad (11)$$

which gives a recipe for computing the edit distance $d(C, \bar{C})$ via dynamic programming. The match cost is found by sequentially updating the dynamic-programming table, and the optimal alignment by tracing through the table [27]. We discretize ψ , as a first approximation, to nine values achieved by using a template, Fig. 3a. This template limits the choices of k and l , which in turn limits the number of segments that can be merged and then matched to three. We have to compute $d(i, j)$ at every point in the 2D grid. Hence, the complexity of matching curve segments is $O(n^2)$, where n is the number of samples along the curve segments. Note that merging does not increase the complexity, as we allow only a fixed number of segments to merge at each point.


Observe that our proposed approach is analogous to the ‘‘edit distance’’ approach [29] for matching strings, where three operations are allowed:

1. relabeling a character,
2. deleting a character, and
3. inserting a character.

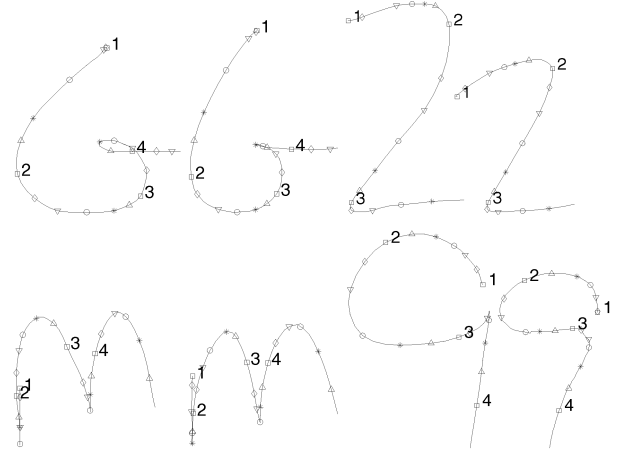
In the curve alignment domain, the infinitesimal curve segment serves the purpose of a ‘‘letter,’’ stretching and bending of infinitesimal segments correspond to the ‘‘relabel edit,’’ removal of an infinitesimal segment of C is analogous to the ‘‘delete edit,’’ and the removal of an infinitesimal segment of \bar{C} is analogous to the ‘‘insert edit.’’ The key difference between string edit distance [29] and our algorithm is that we allow for a merge operation of multiple curve segments. Our algorithm is also similar to the attributed string matching of Tsai and Yu [23]. Other curve matching methods based on dynamic programming [2], [3], [24] also use a similar update step to compute the cost. The main difference in our approach is the use of a fixed template which keeps the complexity at $O(n^2)$.

3.1 Aligning Closed Curves

The main difference between aligning curve segments and closed curves is that, in the former, the start points of the two curves C and \bar{C} correspond, i.e., s_1 is mapped to \bar{s}_1 , while in the latter, the start point correspondence is not known. In order to match closed curves, we need to find the optimal start point correspondence as well. Note that it is sufficient to fix a start point \bar{s}_1 on curve \bar{C} , and then find the optimal alignments for all possible start points on the curve C . It is straightforward to see why. Let ψ_1^* be the optimal alignment when



(a)



(b)

Fig. 6. (a) Sample handwritten characters in our database. (b) Examples of the optimal correspondence. The alignment is indicated by numbers, and within each portion by similar shaped icons. Observe that the matches are intuitive.

we fix \bar{s}_1 as the start point on \bar{C} . Let s_1 be the optimal start point on C and s_2 be the corresponding point of \bar{s}_2 under ψ_1^* . Instead of \bar{s}_1 , if we now choose \bar{s}_2 as the start point on \bar{C} , one alignment maps s_2 to \bar{s}_2 , and s_1 to \bar{s}_1 . This alignment is same as ψ_1^* , and is the optimal alignment due to Proposition 2. Hence, we fix the start point on curve \bar{C} and find optimal alignments for all start points on C .

This may appear to increase the computational complexity by the number of sample points along the closed curve, i.e., to $O(n^3)$, as in previous approaches [24], [3], [2]. Heuristics have been used to reduce the need of having to optimize over n^2 starting point correspondences. For example, Gdalyahu and Weinshall [2] use the distribution of optimal rotation parameters necessary to align the first few points to select plausible start point correspondences. The approaches of Ueda and Suzuki [24] and Milios and Petrakis [3] are not guaranteed to find the optimal alignment for closed curves. As noted in [3], the approach by Ueda and Suzuki [24] may not find the optimal correspondence as only the cost of the best path is stored at each cell in the dynamic-programming table. Milios and Petrakis [3] alleviate this problem by storing the top κ options at each cell, but the optimal correspondence is not guaranteed.

The need to find the optimal alignment for each start point correspondence, in order to ensure that the global minimum of the functional for aligning closed curves is found is rather restrictive, but can be resolved by observing that the optimal alignment curves corresponding to different start points s_i cannot intersect. Specifically, let α_i^* be the optimal alignment curve from $(s_i, 0)$ to $(s_i + L, \bar{L})$, and α_j^* from $(s_j, 0)$ to $(s_j + L, \bar{L})$. Then, for any

TABLE 1
Top 20 Matches for Two Handwritten Digits

2	2 ⁴	2 ⁵	2 ⁸	2 ⁸	2 ⁹	2 ⁹	2 ¹⁰	2 ¹¹	2 ¹¹	2 ¹¹	2 ¹³	2 ¹⁴	2 ¹⁴	7 ¹⁴	7 ¹⁵	7 ¹⁵	7 ¹⁵	7 ¹⁵	3 ¹⁸	
3	3 ⁵	3 ⁶	3 ⁶	3 ⁶	3 ⁷	3 ⁷	3 ⁷	3 ⁷	3 ⁷	3 ⁸	3 ⁸	3 ¹⁰	3 ¹¹	7 ¹⁵	7 ¹⁶	7 ¹⁶	7 ¹⁶	3 ¹⁶	7 ¹⁷	7 ¹⁷

We used 90 handwritten digits (15x{2,3,6,7,8,9}) in this experiment. The number next to the matching character is the Euclidean distance to the query.

TABLE 2
Comparison of Several Methods for Handwritten Character Recognition

Eucl.	0	0 ⁴	0 ⁷	U ¹⁰	6 ¹³	6 ¹⁴	Edit	6	0 ⁶⁹	6 ⁷⁰	0 ⁸¹	6 ⁸²	U ⁹⁴	<table border="1"> <thead> <tr> <th>Method</th><th>Errors</th><th>% recog.</th> </tr> </thead> <tbody> <tr> <td>Euclidean</td><td>4</td><td>98.8%</td> </tr> <tr> <td>Edit</td><td>38</td><td>88.4%</td> </tr> <tr> <td>Hungarian</td><td>52</td><td>84.1%</td> </tr> <tr> <td>ICP</td><td>63</td><td>80.7%</td> </tr> <tr> <td>Softassign</td><td>72</td><td>69.9%</td> </tr> </tbody> </table>	Method	Errors	% recog.	Euclidean	4	98.8%	Edit	38	88.4%	Hungarian	52	84.1%	ICP	63	80.7%	Softassign	72	69.9%
Method	Errors	% recog.																														
Euclidean	4	98.8%																														
Edit	38	88.4%																														
Hungarian	52	84.1%																														
ICP	63	80.7%																														
Softassign	72	69.9%																														
Eucl.	2	2 ⁴	7 ¹¹	2 ¹¹	3 ¹³	3 ¹⁵	Edit	N	2 ⁸²	N ⁹⁶	2 ¹⁰⁶	2 ¹¹⁸	L ¹²⁰																			
Eucl.	6	6 ⁵	6 ¹²	0 ¹⁹	0 ¹⁹	U ¹⁹	Hung.	S	5 ¹⁷³	S ¹⁷⁶	3 ¹⁸¹	2 ¹⁹⁸	3 ²⁰⁵																			
Eucl.	L	L ⁴	1 ¹¹	1 ¹⁵	1 ¹⁶	3 ¹⁸	Hung.	U	V ¹⁵⁶	U ¹⁷⁰	2 ¹⁹⁵	L ²⁰³	2 ²⁰⁸																			
Eucl.	V	V ⁵	U ⁸	0 ¹¹	0 ¹²	P ¹³	ICP	C	0 ⁸⁸	C ⁹²	0 ⁹⁵	P ¹¹¹	2 ¹¹⁸																			
Eucl.	Z	Z ³	2 ⁷	7 ¹⁰	3 ¹³	3 ¹⁴	ICP	I	L ⁴⁰	I ⁵¹	2 ⁵⁵	Z ⁵⁸	W ⁵⁸																			
Eucl.	0	U ¹⁰	0 ¹¹	0 ¹¹	V ¹²	6 ¹⁷	Soft.	S	S ³¹	5 ³³	9 ³⁶	V ⁴⁰	3 ⁶⁷																			
Eucl.	N	W ¹²	N ¹³	U ¹⁹	M ²⁰	- ²²	Soft.	8	3 ⁶⁰	8 ⁶⁸	3 ⁷⁶	2 ⁸⁷	M ⁹²																			

The top five matching categories for a few sample characters using different distance measures are shown. The number next to each match is the cost of the match in each case. The left table shows results of using the Euclidean distance measure. The first six rows show examples where the character is correctly recognized, whereas the last two rows show examples where the top match is incorrect. The table in the middle shows some typical examples where recognition using edit distance, Hungarian assignment based on shape context [31], ICP [32], and Softassign [33] gives incorrect results. Edit distance underperformed in this task mainly due to the difficulties discussed in Section 2.4. Shape context [31] uses global features and has problems differentiating between characters like "U" and "V." Both ICP [32] and softassign [33] allows partial matches, leading to confusion between characters likes "1" and "L." The table on the right summarizes the recognition rates for all the methods.

$s \in [s_i, s_j]$, the optimal alignment curve from $(s, 0)$ to $(s + L, \bar{L})$ has to lie between α_i^* and α_j^* , and this leads to a significant reduction in the search space, Fig. 3b. This is similar to the technique used by Maes [30] for the cyclic string-to-string correction problem. To capture the cyclic nature of the closed curves, the first curve is duplicated in the dynamic-programming grid, Fig. 3b, and we compute the optimal alignment curves

$$\alpha_i^* = (\alpha_{i1}, \dots, \alpha_{ik}), i = 1, \dots, n,$$

where $\alpha_{i1} = (s_i, \bar{s}_1)$, $\alpha_{ik} = (s_{i+n-1}, \bar{s}_m)$, and k is the number of points in α_i^* . When each optimal alignment curve is computed, the search space is split into two, which in analogy to binary search adds a function of $\log(n)$ to the overall complexity, for a total $O(n^2 \log(n))$ for matching closed curves. Fig. 4 plots the computational time for aligning both curve segments and closed curves.

3.2 Effect of Scale

Note that, due to the strict use of intrinsic properties, the proposed curve alignment is invariant to relative rotations and translations of the two curves. However, the dissimilarity measure is not scale-invariant because the stretching term in the functional is not scale-invariant. Methods that rely on using locally scale invariant measures [2] often have difficulty matching curves with the same but largely unequal sized parts. Thus, the determination of scale requires a consideration of the entire curve. This can be done by finding a global optimal scale parameter λ for one curve to match the other and then minimizing over several scales. Since this is asymmetric, we will instead scale up one curve by $\sqrt{\lambda}$ and scale down the other curve by the same factor to reach a common curve, giving rise to the functional

$$\mu_\lambda[g] = \int |\sqrt{\lambda} ds - \frac{1}{\sqrt{\lambda}} d\bar{s}| + R|d\theta - d\bar{\theta}|. \quad (12)$$

The optimal scaling factor λ^{opt} is then computed as $\arg \min_{\lambda} \mu_\lambda[g]$, which can be computed using gradient descent. Note that only a small range of λ , bounded by the ratios of the largest and smallest sampling intervals, needs to be examined, since the functional is strictly increasing beyond this rather limited range.

3.3 Effect of Sampling on Curve Matching

Our approach is robust to the absolute sampling of the curves, i.e., when sampling rates are say doubled or tripled on both curves. However, a differential in the relative sampling rate of aligned curve segments can affect the optimal correspondence and the

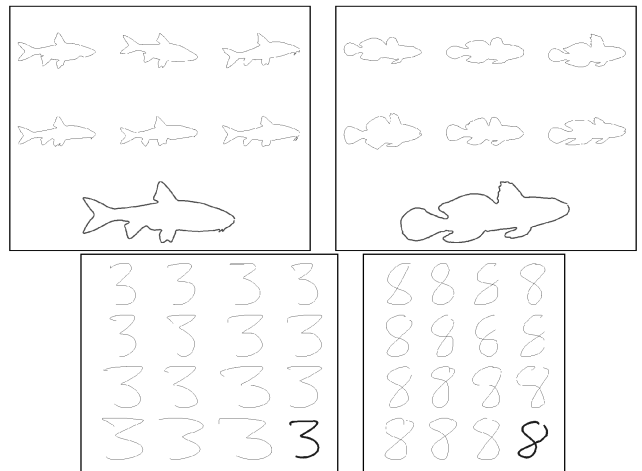


Fig. 7. Original shapes and the computed averages (bold). Observe that in all cases, the average curves are meaningful, and are recognizable as an instance of that category.

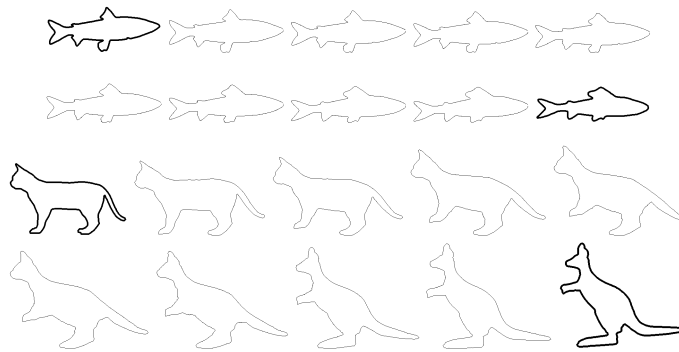


Fig. 8. Two examples of deforming a shape outline to another. The original shapes are shown in bold. Observe that the intermediate shapes are meaningful.

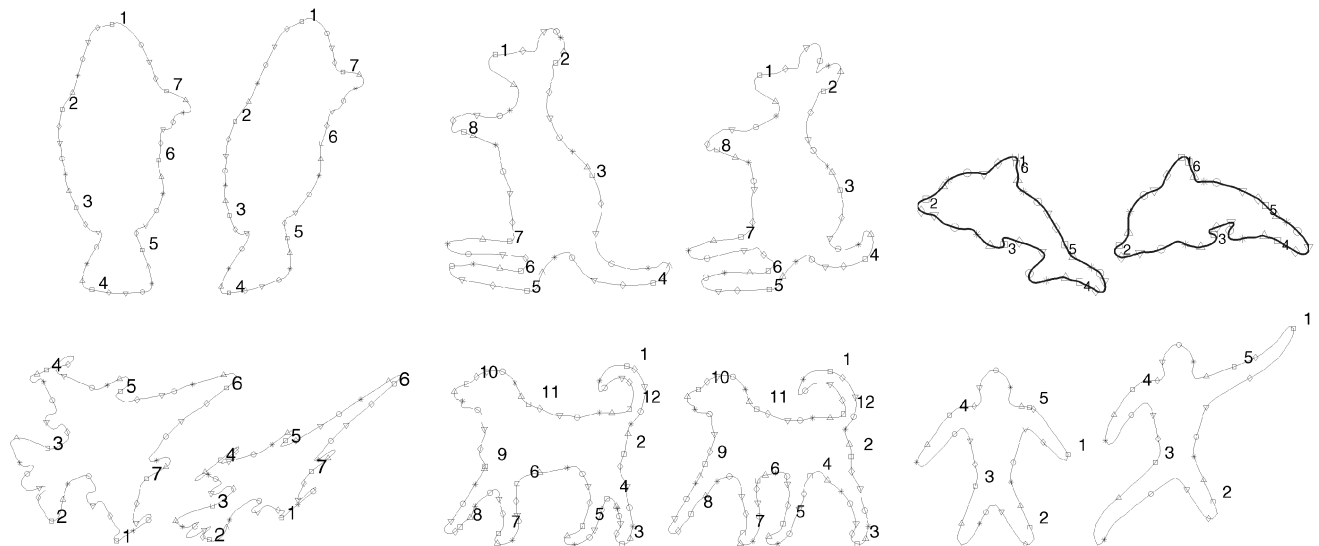


Fig. 9. This figure illustrates the performance of curve alignment in the presence of **affine transformation** (left column), **viewpoint variation** (middle column), and **articulation and deformation of parts** (right column). Curve matching gives the intuitive correspondence in all cases, i.e., salient parts correspond.

match cost. Fig. 5 illustrates this effect in two cases. In the first case, Fig. 5a, the sampling rates are only somewhat different, the correspondence is good and the cost is affected only slightly. The correspondence and, hence, the cost can be improved by a postprocessing step by “fine-tuning” the computed alignment curve points by allowing them to move along the horizontal and vertical grid-lines to minimize the match cost, Fig. 5c. In the second case, when the sampling rate is substantially different along corresponding curve segments, the alignment and costs are affected more significantly, Fig. 5b. This effect is best handled by resorting to larger templates, requiring additional computational effort. As a practical alternative the curves can be resampled a priori, to ensure that the sampling rates are not drastically different.

4 APPLICATIONS OF CURVE ALIGNMENT

In this section, we illustrate that the proposed curve alignment is fairly generic and useful across many applications.

4.1 Handwritten Character Recognition

Handwritten character data is typically acquired by sampling the trace of the writer’s pen and is represented by a pair of coordinates (x, y) , resulting in a curve segment, Fig. 6a. Handwritten character data is thus inherently one-dimensional in nature, rendering curve alignment as a natural choice to measure character similarity. Fig. 6b shows a few examples of the alignment found between two

characters using our framework. We use a database of six different digits with 15 samples of each for a total of 90 items. Matching is done between each pair of digits in the database. The top 20 matches of a few sample digits are shown in Table 1. Note that in most cases the top matches are other samples of the same digit (one error only). This is a typical result.

Working with a large database necessitates the use of character prototypes [6], [5] to reduce the number of matches required. The prototype can either be a representative sample [5] or an average curve [6]. We use an “average” character as the prototype, Section 4.2. The cost of matching each “unknown” character to all the average characters is computed and the nearest neighbor is chosen as the correct match. We used a database of 327 digits and alphabets (34 categories) written by one subject. The top five matching prototypes for a few sample characters are shown in Table 2. We consider the character to be correctly recognized if the top matching prototype is correct, i.e., we say a “2” is recognized correctly if the top prototype is “2.” Out of 327 characters, 323 were correctly recognized using the Euclidean measure, leading to a recognition rate of 98.8 percent. The performance of other methods was significantly lower, see Table 2.

4.2 Curve Averaging and Morphing

The ability to find the average of a set of curves has several applications in computer vision including prototype formation, shape morphing, and computational atlases. We now show that using the curve alignment framework, the average curve can be computed by averaging the intrinsic properties of the corresponding

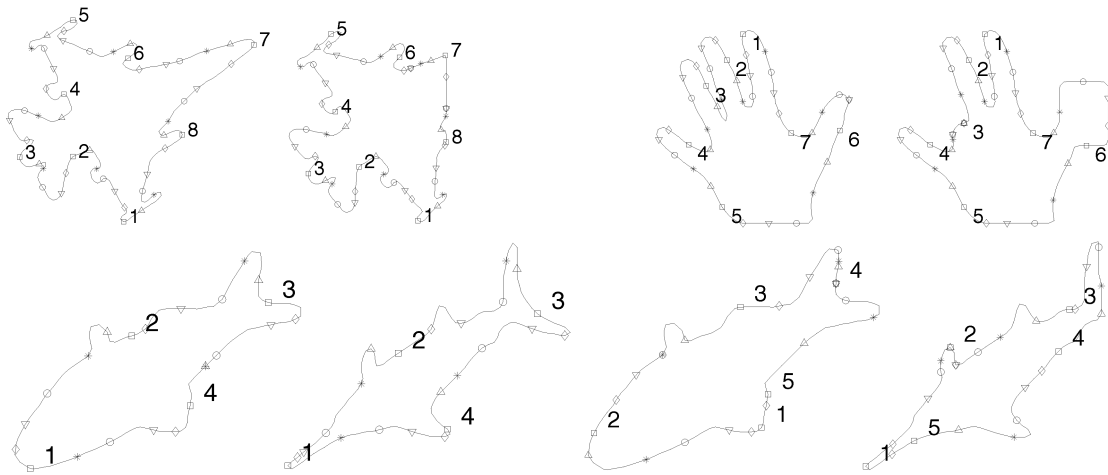


Fig. 10. This figure illustrates the performance of curve alignment in the presence of **partial occlusion**. Top row shows two examples where curve matching gives the intuitive correspondence. Observe that the effects of occlusion is limited to the affected parts. Bottom row illustrates that curve alignment may give the unintuitive correspondence in presence of occlusion that affects the overall part structure. Bottom left: The correspondence between the fishes is intuitive, with the tails, heads, and the fins correctly matched. Bottom right: Observe the effect of occluding part of the tail of the fish on the right. The occlusion, though small, affects the part structure of the shape, and gives the wrong correspondence, with the head of the fish on the left matched to a fin on the fish on the right.

curve subsegments. Fig. 7 shows the average curve for fish outlines and handwritten digits. Observe that in each case the average curve maintains the general shape as well as the local features and their relationships, and is recognizable as an instance of that category. Note that the average handwritten characters were used as prototypes in the handwritten character recognition experiments of Section 4.1 with excellent results.

The proposed curve alignment framework can also be used to generate a morphing sequence to “interpolate” between two shapes that are not very dissimilar. The main idea is to first find the optimal correspondence between the two shapes, and then deform each curve segment to its corresponding segment, using a weighted average of their intrinsic properties. Fig. 8 shows examples of morphing one type of fish to another, and a cat to a kangaroo. Observe that all the intermediate shapes are meaningful in these examples, although in some examples self intersections can occur. This approach is only appropriate for somewhat similar shapes, and the general case requires a notion of the interior, e.g., as provided by the medial axis.

4.3 Object Recognition

We now examine the effectiveness of curve alignment for matching shapes which are represented by their outlines. For curve alignment to be effective in object recognition, it has to perform well under a variety of visual transformations. Fig. 9 shows that the curve matching algorithm works well in the presence of some of the commonly occurring visual transformations like affine transformations, modest amounts of viewpoint variation, articulation, stretching, and bending of parts. Fig. 10 examines the effects of partial occlusion. Curve matching gives the intuitive correspondence in the presence of occlusion, when the spatial arrangement of parts is not significantly affected by the occlusion, Fig. 10 (top

row). On the other hand, it may fail when occlusion affects the overall part structure of the shape, Fig. 10 (bottom row).

We examine the effectiveness of the proposed edit distance between shape outlines on the database of shapes created for testing the compression rates for MPEG7 [34]. The database consists of 70 categories with 20 shapes in each category for a total of 1,400 shapes. The performance is measured by counting the number of correct matches in the top 40 matches [34]. Because our approach is not invariant to the “flip” transformation, for each pair of shapes, we compute the optimal alignments for the shapes before and after flipping one of the shapes, and take the minimum of the two as the match cost. The retrieval rate for our approach is 78.17 percent. The currently published best performing approaches use curvature scale space [35], correspondence of visual parts [34], and matching shape contexts [36], with retrieval rates of 75.44 percent, 76.45 percent, and 76.51 percent, respectively, see Table 3.

Finally, we have applied our curve alignment method to the GESTURE data set created by Milios and Petrakis [3], who also provided us with their results and human similarity judgements data. In their paper, they evaluated four methods, namely, Fourier descriptors [37], sequential moments [38], geometric moments [39], and their method based on dynamic programming [3]. Human similarity judgements on the top 50 matches were used to judge whether a match is correct and evaluated using a precision-recall diagram.⁴ We augmented these results by applying shape contexts using the Hungarian method [31] and curve alignment to this data set. The results for all methods are plotted using the precision-recall diagram, Fig. 11. Observe that the proposed method significantly outperforms all the competing methods.

A key advantage of the proposed curve alignment for recognition tasks is that it can be used in practical applications due to its high efficiency and excellent recognition rates, even under a moderate range of visual transformations. The main drawback of using this approach when applied to recognition of shapes is that the curve-based representation does not explicitly capture the regional aspect of a shape and is sensitive to the presence and spatial arrangement of parts, Fig. 10. Nevertheless, this framework is useful for a variety of applications involving indexing into and organizing a database of shapes.

4. Precision is the percentage of similar shapes retrieved with respect to total number of retrieved shapes. Recall is the percentage of similar shapes retrieved with respect to total number of similar shapes in the database.

TABLE 3
The Retrieval Rates for MPEG7 Shape Database

Method	Performance
CSS [35]	75.44%
Visual Parts [34]	76.45%
Shape Contexts [36]	76.51%
Curve Edit Distance	78.17%

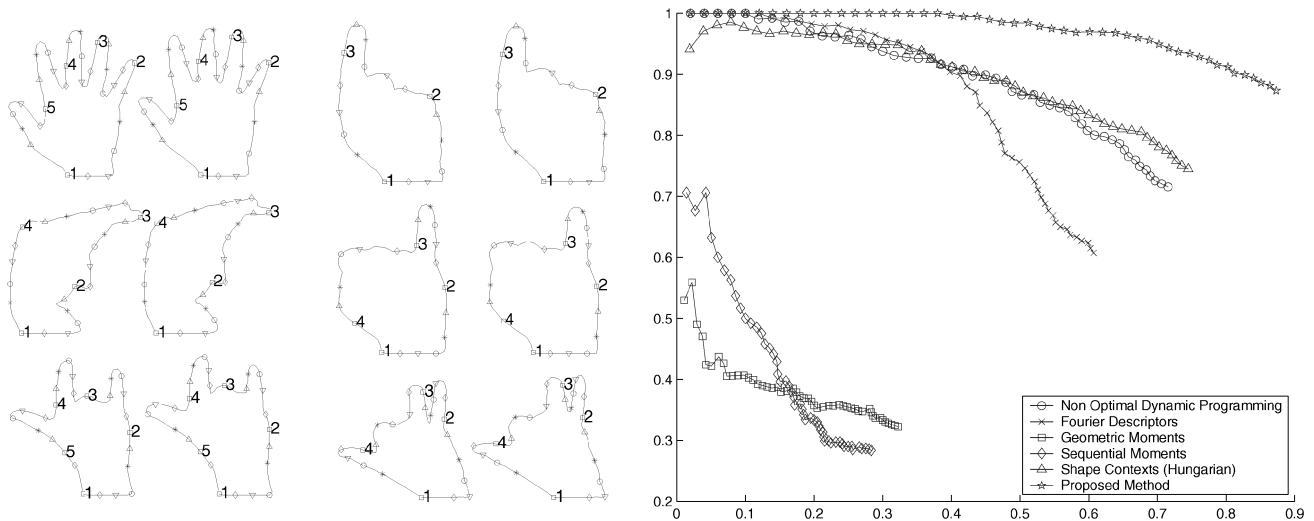


Fig. 11. Left: The result of matching some representative gestures. Observe that the correspondences is intuitive in all cases. Right: Precision-recall diagram for the GESTURE data set. The results of Fourier descriptors, sequential moments, geometric moments, and nonoptimal dynamic programming approaches are from the original paper. We have added two additional results, shape contexts and our technique to this data set. Observe that the proposed approach outperforms all the competing methods.

5 CONCLUSION

We have presented a novel method for finding the optimal correspondence between two curves based on the concept of an *alignment curve*. The optimal alignment is then used to define a metric of dissimilarity between two shapes. We have developed a fast dynamic-programming algorithm for finding the optimal alignment between pairs of curve segments and also pairs of closed curves. The exclusive use of intrinsic properties in our formulation leads to invariance under rigid transformation. Our experiments show that curve alignment is robust under a variety of visual transformations. We have illustrated the usefulness of the alignment curve approach for a variety of computer vision applications including handwritten character recognition, prototype formation, shape morphing, and object recognition.

ACKNOWLEDGMENTS

B. Kimia acknowledges the support of the US National Science Foundation (NSF) grants IRI-0083231 and BCS-9980091. P. Klein acknowledges the support of NSF Grant CCR-9700146. The authors are grateful to Farzin Mokhtarian for providing the fish outlines and Jayashree Subrohmnia at IBM Pen Technologies group for providing the handwritten character data.

REFERENCES

- [1] N.J. Ayache and O.D. Faugeras, "HYPER: A New Approach for the Recognition and Positioning of Two-Dimensional Objects," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 8, no. 1, pp. 44-54, 1986.
- [2] 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*, vol. 21, no. 12, pp. 1312-1328, 1999.
- [3] E. Milios and E.G.M. Petrakis, "Shape Retrieval Based on Dynamic Programming," *IEEE Trans. Image Processing*, vol. 9, no. 1, pp. 141-146, 2000.
- [4] C.C. Tappert, "Cursive Script Recognition by Elastic Matching," *IBM J. Research Development*, vol. 26, no. 6, pp. 765-771, 1982.
- [5] S. Connell and A.K. Jain, "Learning Prototypes for On-Line Handwritten Digits," *Proc. Int'l Conf. Pattern Recognition*, pp. 182-184, 1998.
- [6] B. Wirtz, "Average Prototypes for Stroke-Based Signature Verification," *Proc. Int'l Conf. Document Analysis and Recognition*, pp. 268-272, 1997.
- [7] I. Cohen, N. Ayache, and P. Sulger, "Tracking Points on Deformable Objects Using Curvature Information," *Proc. European Conf. Computer Vision*, pp. 458-466, 1992.
- [8] H.D. Tagare, "Shape-Based Nonrigid Correspondence with Application to Heart Motion Analysis," *IEEE Trans. Medical Imaging*, vol. 18, no. 7, pp. 570-578, 1999.
- [9] J.T. Schwartz and M. Sharir, "Identification of Partially Obscured Objects in Two and Three Dimensions by Matching Noisy Characteristic Curves," *Int'l J. Robotics Research*, vol. 6, no. 2, pp. 29-44, 1987.
- [10] S. Umeyama, "Parameterized Point Pattern Matching and its Application to Recognition of Object Families," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 15, no. 2, pp. 136-144, 1993.
- [11] R. Basri, L. Costa, D. Geiger, and D. Jacobs, "Determining the Similarity of Deformable Shapes," *Vision Research*, vol. 38, pp. 2365-2385, 1998.
- [12] L. Younes, "Computable Elastic Distance between Shapes," *SIAM J. Applied Math.*, vol. 58, pp. 565-586, 1998.
- [13] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active Contour Models," *Int'l J. Computer Vision*, vol. 1, no. 4, pp. 321-331, 1988.
- [14] D.W. Jacobs, D. Weinshall, and Y. Gdalyahu, "Classification with Nonmetric Distances: Image Retrieval and Class Representation," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 22, no. 6, pp. 583-600, 2000.
- [15] A. Tversky, "Features of Similarity," *Psychological Rev.*, vol. 84, no. 4, pp. 327-352, 1977.
- [16] T.B. Sebastian, P.N. Klein, and B.B. Kimia, "Shock-Based Indexing into Large Shape Databases," *Proc. Seventh European Conf. Computer Vision*, May 2002.
- [17] S. Brin, "Near Neighbor Search in Large Metric Spaces," *Proc. Int'l. Conf. Very Large Databases (VLDB)*, pp. 574-584, 1995.
- [18] J. Uhlmann, "Satisfying General Proximity/Similarity Queries with Metric Trees," *Information Processing Letters*, vol. 40, pp. 175-179, 1991.
- [19] P. Yianilos, "Data Structures and Algorithms for Nearest Neighbor Search in General Metric Spaces," *ACM-SIAM Symp. Discrete Algorithms*, pp. 311-321, 1993.
- [20] A. Marzal and E. Vidal, "Computation of Normalized Edit Distances and Applications," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 15, pp. 926-932, 1993.
- [21] K.S. Arun, T.S. Huang, and S.D. Blostein, "Least-Squares Fitting of Two 3-D Point Sets," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 9, no. 5, pp. 698-700, 1987.
- [22] H.C. Liu and M.D. Srinath, "Partial Shape Classification Using Contour Matching in Distance Transformation," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 12, no. 11, pp. 1072-1079, 1990.
- [23] W.H. Tsai and S.S. Yu, "Attributed String Matching with Merging for Shape Recognition," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 7, no. 4, pp. 453-462, 1985.
- [24] N. Ueda and S. Suzuki, "Learning Visual Models from Shape Contours Using Multiscale Convex/Concave Structure Matching," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 15, no. 4, pp. 337-352, 1993.
- [25] H.D. Tagare, "Deformable 2-D Template Matching Using Orthogonal Curves," *IEEE Trans. Medical Imaging*, vol. 16, no. 1, pp. 108-117, 1997.
- [26] H. Sakoe and S. Chiba, "Dynamic Programming Algorithm Optimization for Spoken Word Recognition," *IEEE Trans. Acoustics, Speech, and Signal Processing*, vol. 26, no. 1 pp. 43-49, 1978.
- [27] R.E. Bellman and S.E. Dreyfus, *Applied Dynamic Programming*. Princeton Univ. Press, 1962.
- [28] T. Sebastian, P. Klein, and B. Kimia, "Curve Matching Using Alignment Curve," Technical Report LEMS 184, LEMS, Brown Univ., 2000.
- [29] R.A. Wagner and M.J. Fischer, "The String-to-String Correction Problem," *J. ACM*, vol. 21, pp. 168-173, 1974.