

Extensions of Invariant Signatures for Object Recognition

Daniel J. Hoff¹
Department of Mathematics
University of California, San Diego
La Jolla, CA 92093
d1hoff@math.ucsd.edu

Peter J. Olver¹
School of Mathematics
University of Minnesota
Minneapolis, MN 55455
olver@math.umn.edu
<http://www.math.umn.edu/~olver>

Abstract

A refinement of the method of differential invariant signatures for object recognition is presented. The value of the method lies in its compromise between local and global identifying properties, thereby allowing us to distinguish non-congruent curves whose Euclidean signatures have identical trace.

Keywords:

object recognition, plane curve, rigid equivalence, curvature, vertex, bivertex arc, differential invariant, Euclidean signature

1 Introduction.

One of the primary goals of modern image processing is to recognize objects in different orientations and positions. In simple situations, this translates to an analysis of the object's boundary under various geometric transformations. In this paper, we concentrate on planar images, deferring the more challenging cases of three-dimensional images and videos to subsequent investigations. Objects are represented by their boundary curves $C \subset \mathbb{R}^2$. Given a transformation group G acting on \mathbb{R}^2 , our task is to determine whether or not two objects can be mapped to each other by a transformation $g \in G$, or, in other words, whether the boundary curves C, \tilde{C} are *congruent* under G , meaning that $\tilde{C} = g \cdot C$ for some $g \in G$. If $\tilde{C} = C$, then the congruence problem reduces to the classification of *symmetries* of the boundary curve, meaning group elements that map C to itself. The most important non-trivial case, and the subject of this paper, takes $G = \text{SE}(2)$ to be the special Euclidean group of rigid planar motions: rotations and translations. Other examples relevant to image processing include the special affine group of area-preserving affine transformations and the projective group; the methods we develop here can be systematically extended to these and many additional contexts.

In [2], inspired by Élie Cartan's solution to equivalence problems for submanifolds, [3, 11], and the equivariant method of moving frames, [4], the authors proposed a general theory of G -invariant signatures for use in object recognition and symmetry detection. The signature submanifold is parametrized by a distinguished, finite collection of fundamental differential

¹Supported in part by NSF Grant DMS 08-07317.

invariants associated with the transformation group G . Under appropriate regularity conditions, [11], two submanifolds are locally congruent under G if and only if their signatures are identical. In the case $G = \text{SE}(2)$ acting on plane curves, there are two fundamental differential invariants: the curvature, κ , and its first derivative with respect to arc length, κ_s , which together parametrize the Euclidean signature curve $S = \{(\kappa, \kappa_s)\}$. In general, differential invariant signatures offer several important advantages over other proposed object recognition methods. First, they are completely local, in that the differential invariants parametrizing the signature are all measured at a single point of the object. Thus, occlusions and comparison of parts of objects are readily incorporated into the framework since the partial object delineates a well-defined subset of the complete signature. Boundary recognition methods that depend explicitly on global properties such as arc-length are not nearly as robust under such a loss of information. Second, differential invariant signatures can be readily used to recognize symmetries and approximate symmetries of the object, since the signature's *index*, meaning the number of points in the original submanifold that map to a generic signature point, is precisely the number of discrete symmetries contained within the specified transformation group. Further, the subsequent equivariant method of moving frames, [4], provides a simple, algorithmic means for constructing differential invariant signatures for curves, surfaces, and higher dimensional submanifolds under a very broad range of transformation groups that includes all examples of relevance to image processing. While the presence of high derivatives in the signature invariants means that noise sensitivity will be a concern, our implementations and applications reconfirm the robustness of the signature approach, provided one incorporates a preliminary smoothing of the object boundary followed by suitable invariant numerical schemes to approximate the curvature invariants.

The main theorem governing the congruence of submanifolds with identical signatures was, unfortunately, stated in [2] without explicit mention of the regularity conditions required for its validity. As a result, readers not familiar with the underlying mathematics have tried applying it to irregular curves and surfaces, for which its validity is no longer guaranteed. Indeed, Musso and Nicolodi, [9], recently constructed one-parameter families of smooth simple closed curves that are not rigidly congruent and yet their Euclidean signatures have identical trace. Their examples served to motivate our constructions, in which we augment the signature comparison method in a way that can be applied to a wider class of boundary curves. Specifically, we show that, by properly partitioning the boundary curves into subarcs, a combination of signature comparison and transformation reconstruction is able to classify the curves up to rigid motion, even in the presence of noise.

Our approach is, in part, inspired by the methods developed in [12], in which contour shape is encoded by *codons*, which are touted as “primitive part descriptors”. In the case of plane curves under rigid motion, the codons used in [12] are subarcs whose endpoints are successive minima of curvature. We refine this idea by employing *bivertex arcs*, whose endpoints are successive stationary points of curvature, that is, where its derivative with respect to arc length vanishes. Any reasonable — meaning “v-regular” as per Definition 1 — non-circular Jordan curve can be decomposed into a non-overlapping union of finitely many bivertex arcs together with, possibly, a finite number of circular arcs and straight line segments. Our comparison algorithm concentrates exclusively on matching bivertex arcs up to rigid motion, using their Euclidean signatures, and then determining whether the resulting rigid motions are sufficiently close to each other to produce a global congruence of the two

boundary curves. As such, our approach incorporates all the advantages of parts-based methods of object representation, [8, 12], including the ability to match pieces of objects as well as reconstruct objects under partial occlusion. In this vein, further refinements, supported by applications to the automatic assembly of jigsaw puzzles can be found in [7].

It is important to keep in mind that, in practice, objects segmented from digital images are actually represented as discrete point sets that approximately sample their boundary curves. Euclidean-invariant numerical algorithms for approximating signature curves from discrete data have been developed in [2], and then improved by Boutin, [1]. Thus, in practical implementations, rather than work directly with the signature curves, we base our comparison algorithms on their discretizations. The discretized curves are then decomposed into (approximate) discrete bivertex arcs. Comparison of pairs of bivertex arc signatures will be based on a generalized electrostatic attraction between their discretizations, viewed, respectively, as sets of positive and negative point charges. Full details can be found within the body of the paper.

2 Generalized Euclidean Signatures.

Let us introduce our basic terminology and assumptions. We will consider simple (non-intersecting) plane curves $C \subset \mathbb{R}^2$ of class at least C^3 . Let $z(s)$, $0 \leq s \leq L$, denote the arc length parametrization, so that L is the length of C . Our primary focus is on closed Jordan curves, for which $z(s)$ is periodic of period L . In contrast, a simple curve with distinct endpoints will be called a (closed) *arc*. An *open arc* is obtained by removing the endpoints.

Let $\kappa(s) = z'(s) \wedge z''(s)$ denote the signed *curvature* at the point $z(s) \in C$, with \wedge denoting the scalar cross product: $v \wedge w = v_1 w_2 - v_2 w_1$ for $v = (v_1, v_2)$, $w = (w_1, w_2) \in \mathbb{R}^2$. We use $\kappa_s = d\kappa/ds$ to denote the derivative of curvature with respect to the arc length parameter. Both κ and κ_s (as well as all higher order arc length derivatives) are Euclidean *differential invariants*, meaning that they are unchanged under rigid motion, [11]. A point $z(s) \in C$ is called *regular* if $\kappa_s(s) \neq 0$. An *ordinary vertex* is a local extremum of curvature, [6]. By a *generalized vertex*, we mean a maximal connected, closed arc $V \subset C$ on which $\kappa_s(s) \equiv 0$. Thus, a generalized vertex is either an ordinary vertex, or a critical point of curvature, or a circular arc, or a straight line segment. For us, the term “vertex” will always mean “generalized vertex” — in contrast to most works, where it means “ordinary vertex”. We let $C_v = \{\kappa_s = 0\} \subset C$ denote the union of all vertices, while the rest, $C_r = \{\kappa_s \neq 0\} = C \setminus C_v$, forms the *regular part* of the curve.

Definition 1. A C^3 curve is called *v-regular* if it has finitely many generalized vertices.

The requirement of v-regularity excludes pathological smooth curves that have infinitely many vertices. Such curves are merely mathematical curiosities, and do not play any role in digital imagery. From now on, the Jordan curve C is assumed to be v-regular, and so $C_v = V_1 \cup \dots \cup V_m$ is a finite union of points, circular arcs, and straight line segments. Unless the curve is a circle, its regular part $C_r = A_1 \cup \dots \cup A_m$ is a finite union of $m \geq 4$ disjoint open arcs. (The lower bound on m follows from a straightforward adaptation of Mukhopadhyaya’s Four Vertex Theorem, following the proof in [6, section 2.3].) We let $B_j = \overline{A_j}$ denote the corresponding closed arcs, which we distinguish with a special name:

Definition 2. A *bivertex arc* is a simple arc, of class C^3 , whose endpoints belong to vertices, so $\kappa_s = 0$ at both ends, but all of whose interior points are regular, so that $\kappa_s \neq 0$ away from the endpoints.

Observe that curvature κ varies monotonically along a bivertex arc, which is thus a special kind of spiral arc, [6].

Lemma 3. *Any v -regular simple closed curve that is not a circle can be uniquely decomposed,*

$$C = \bigcup_{j=1}^m B_j \cup \bigcup_{k=1}^n V_k, \quad (2.1)$$

into a finite union of $m \geq 4$ non-overlapping¹ bivertex arcs B_1, \dots, B_m , and $n \geq 0$ generalized vertices V_1, \dots, V_n , each a circular arc or straight line segment.

We will refer to (2.1) as the *bivertex decomposition* of the curve C . Note that we exclude point vertices from the bivertex decomposition, since they are accounted for by the endpoints of the bivertex arcs B_j .

Two plane curves $C, \tilde{C} \in \mathbb{R}^2$ are said to be *rigidly equivalent*, or *congruent* for short, if there exists a rigid motion $g \in \text{SE}(2)$ such that $\tilde{C} = g \cdot C$. We extend the notion of congruence to disconnected unions of arcs, keeping in mind that for two unions to be congruent, their constituent vertices and arcs must be pair-wise congruent *under the same rigid motion*.

Theorem 4. *Two v -regular non-circular Jordan curves $C, \tilde{C} \subset \mathbb{R}^2$ are congruent if and only if their bivertex decompositions contain the same number of bivertex arcs and generalized vertices, and, moreover, there exists a common rigid motion $g \in \text{SE}(2)$ such that $\tilde{B}_j = g \cdot B_j$, $j = 1, \dots, m$, and $\tilde{V}_k = g \cdot V_k$, $k = 1, \dots, n$.*

In fact, in view of the following result, one only needs to consider the regular, bivertex arcs for Theorem 4 to be valid.

Theorem 5. *Two v -regular plane curves $C, \tilde{C} \subset \mathbb{R}^2$ are congruent if and only if their regular parts C_r, \tilde{C}_r are congruent.*

Proof: Since κ_s is a Euclidean differential invariant, (generalized) vertices are preserved under rigid motion, and hence congruence of curves immediately implies congruence of their regular parts. Conversely, suppose $\tilde{C}_r = g \cdot C_r$ for some $g \in \text{SE}(2)$. Let $V \subset C_v$ be a generalized vertex of C . If V is a single point, then it belongs to the closure of C_r , and hence, by continuity, $\tilde{V} = g \cdot V$ is also a point vertex of \tilde{C} . On the other hand, if V is a circular arc or straight line segment, its endpoints z_i , $i = 1, 2$, are also endpoints of bivertex arcs B_i . We claim that $\tilde{V} = g \cdot V$ is a vertex of \tilde{C} whose endpoints $\tilde{z}_i = g \cdot z_i$ are the endpoints of the bivertex arcs $\tilde{B}_i = g \cdot B_i$.

Since C is assumed to be at least C^2 , the curvature of V , i.e., the reciprocal of its radius, is uniquely determined by its value at one endpoint, say z_1 . Moreover, since C is simple, the second endpoint z_2 is *uniquely* characterized as the first point belonging to the closure of C_r that is encountered when tracing the circle/line with the prescribed curvature emanating from z_1 . Again, by continuity, this observation suffices to prove that $\tilde{V} = g \cdot V$ is a generalized vertex connecting the corresponding regular arcs of \tilde{C} . Q.E.D.

¹“Non-overlapping” means they have at most one endpoint in common.

The following definition was proposed in [2].

Definition 6. Let C be a plane curve of class C^3 and of finite length $L < \infty$. The *Eulerian signature* of C is the (non-simple) plane curve $S(C) = \{(\kappa(s), \kappa_s(s)) \mid 0 \leq s \leq L\}$ parametrized by the curvature and its derivative with respect to arc-length.

The next result is a consequence of very general results on group-invariant signatures of regular submanifolds, [4, Theorem 14.9], which is a reformulation of Cartan’s general solution to the local equivalence problem for submanifolds under Lie group actions², [3, 11].

Theorem 7. *Suppose that C and \tilde{C} are simple curves of class at least C^3 that contain no vertices, so $C = C_r$, $\tilde{C} = \tilde{C}_r$. (In particular, neither curve can be closed.) Then the curves are rigidly equivalent, so $\tilde{C} = g \cdot C$, for some $g \in \text{SE}(2)$, if and only if their signatures are identical: $S(C) = S(\tilde{C})$.*

By continuity, Theorem 7 easily extends to bivertex arcs. Moreover, if C is a bivertex arc, then κ_s remains of one sign, and hence its signature $S(C)$ is a simple arc in the (κ, κ_s) -plane, which, by the monotonicity of κ , is described by the graph of a single-valued function $\kappa_s = H(\kappa)$ that starts and ends on the κ axis. On the other hand, the signature of any circular arc or straight line segment is a single point on the κ axis, and hence, by itself, is unable to distinguish between those with different lengths. Indeed, the curves constructed in [9] are exactly of this type, where circular arcs are inserted into a suitable closed curve at its vertices, producing non-congruent smooth Jordan curves that nevertheless all have a common signature trace. The point is that the modified curves no longer satisfy the regularity conditions required for the Cartan equivalence result to be valid. On the other hand, Theorem 7 remains valid as stated if the curve is allowed to contain non-degenerate vertices, meaning those at which $\kappa_{ss} \neq 0$, and thereby retain the regularity requirement.

The following consequence of Theorems 4 and 7 forms the basis for our proposed signature-based method for recognition of rigidly equivalent objects.

Theorem 8. *Let $C, \tilde{C} \subset \mathbb{R}^2$ be v -regular non-circular Jordan curves. Assume that their bivertex decompositions (2.1) contain the same number, m , of non-overlapping bivertex arcs, and, moreover, for each $j = 1, \dots, m$, the bivertex arcs B_j and \tilde{B}_j have identical signatures: $S(B_j) = S(\tilde{B}_j)$, which, by Theorem 7, implies that there exist rigid motions $g_j \in \text{SE}(2)$ such that $\tilde{B}_j = g_j \cdot B_j$, $j = 1, \dots, m$. If, in addition, all the $g_j = g \in \text{SE}(2)$ are the same, then $\tilde{C} = g \cdot C$ are congruent.*

Proof: The existence of each g_j follows immediately from Theorem 7. If, furthermore, all $g_j = g$, then

$$\tilde{C}_r = \bigcup_j \tilde{A}_j = \bigcup_j g \cdot A_j = g \cdot \bigcup_j A_j = g \cdot C_r,$$

where \tilde{A}_j, A_j denote the open, regular arcs obtained by deleting the endpoints of the corresponding bivertex arcs \tilde{B}_j, B_j , respectively. The result now follows from Theorem 5. *Q.E.D.*

²In [11], the older term “classifying submanifold” is used in place of “signature submanifold”.

Remark: The upshot of the preceding discussion is that one should not, in general, regard the trace of the signature curve, $S(C) \subset \mathbb{R}^2$, as a “complete” curve signature, in that it fails to distinguish inequivalent (albeit irregular) curves. A more refined theory could be founded on some sort of appropriately weighted signature curve, where the weight is determined by the local measure of arc length of the original curve. In particular, the signature of a circular arc would be a single point on the κ axis weighted by its overall length; this refinement would allow us to distinguish inequivalent arcs.

Indeed, if we were to discretize the curve C by choosing a reasonably dense set of sample points, uniformly distributed with respect to arc length, then the corresponding discrete signature will consist of (approximations to) the curvature and its derivative at the sample points. Any circular arc or straight line segment contained in C would correspond, not to a single point on the discretized signature, but to several (nearly) identical signature points on or near the κ axis, their number being governed by the length of the arc or segment. The weighted signature for such a curve should, then, reflect this fact, and be in a certain sense, the appropriate limiting object as the sample points become more and more dense. Noise could then be readily incorporated into the weighted signature scheme by use of an appropriate probability distribution centered on the underlying objects. Details of this approach, and its extension to more complicated group actions as well as higher dimensional submanifolds will be developed elsewhere.

3 A Practical Algorithm for Curve Classification.

In this section, we apply the results of Section 2 to devise a practical algorithm for curve comparison and classification. Some preliminary applications will be described in the following section. Further extensions, along with a significant new application to the automatic assembly of jigsaw puzzles, are the topics of [7].

Let C and \tilde{C} be simple, closed, v -regular plane curves, representing the boundaries of two planar objects. In the case of noisy curve data, we can apply appropriate smoothing techniques, e.g., the Euclidean or affine-invariant curve shortening flows, [10], or some form of spline-based smoothing, [13], before initiating the analysis. However, as we observe in the applications, the method is quite robust, even in the presence of noisy data.

We will devise a *congruence coefficient*, $\mu(C, \tilde{C}) \in [0, 1]$, with the convention that a μ score of 1 indicates rigid equivalence of the two curves, while μ decreases to 0 as the curves become “less and less” congruent. The key steps in our comparison method, which is based on Theorem 8, are described in the following subsections:

- 3.1. Decompose C and \tilde{C} into a collection of (approximate) bivertex arcs.
- 3.2. Compare the Euclidean signatures of the individual arcs.
- 3.3. If the signatures match, reconstruct and then compare the rigid motions relating each pair of congruent arcs.

Now, in practical applications to objects in digital images, we only know a discrete set of points $C^\Delta = \{z^1, \dots, z^N\}$, where each z^j lies on or near the actual boundary curve C . The curve can be reconstructed by some form of interpolation, e.g., periodic splines, possibly refined by smoothing to reduce the effect of noise. In our applications, we work directly with the discretized version. We similarly (approximately) discretize the signature curve

S by a set of *signature points* $S^\Delta = \{\sigma^1, \dots, \sigma^N\}$, where each $\sigma^j = (\kappa^j, \kappa_s^j)$ consists of suitable approximations to the curvature and its length derivative at the corresponding sample point $z^j = (x^j, y^j) \in C^\Delta$. For example, the entries of σ^j may be found directly from the discretized curve C^Δ by use of the Euclidean-invariant numerical approximations to the curvature invariants developed in [1, 2].

3.1 Approximating Bivertex Decompositions

We begin by constructing the bivertex decompositions of our two (discretized) curves C and \tilde{C} . In practical situations, we must avoid bivertex arcs that are the result of ambient noise. This is accomplished by selecting some small $\delta_0 > 0$, and splitting each curve into arcs whenever $|\kappa_s| - \delta_0$ changes sign. The collection of connected arcs that satisfy either $\kappa_s > \delta_0$ or $\kappa_s < -\delta_0$ away from their endpoints will represent the bivertex arcs in the bivertex decomposition (2.1) of the curve. We then eliminate insignificant arcs over which curvature changes by less than some prescribed $\delta_1 > 0$. (Keep in mind that κ is monotone on bivertex arcs, and hence the change is merely the absolute value of the difference in its values at the endpoints.) To specify the cut-offs δ_0, δ_1 , let us define the *scale of comparison* of the two curves to be

$$D(C, \tilde{C}) = (D_\kappa(C, \tilde{C}), D_{\kappa_s}(C, \tilde{C})), \quad (3.1)$$

where

$$\begin{aligned} D_\kappa(C, \tilde{C}) &= \max \left\{ \max_{z \in C}(\kappa|_z) - \min_{z \in C}(\kappa|_z), \max_{\tilde{z} \in \tilde{C}}(\kappa|\tilde{z}) - \min_{\tilde{z} \in \tilde{C}}(\kappa|\tilde{z}) \right\}, \\ D_{\kappa_s}(C, \tilde{C}) &= \max \left\{ \max_{z \in C}(\kappa_s|_z) - \min_{z \in C}(\kappa_s|_z), \max_{\tilde{z} \in \tilde{C}}(\kappa_s|\tilde{z}) - \min_{\tilde{z} \in \tilde{C}}(\kappa_s|\tilde{z}) \right\}, \end{aligned} \quad (3.2)$$

represent the maximal variation of, respectively, κ and κ_s along the two curves. We then set

$$\delta_0 = \frac{D_{\kappa_s}(C, \tilde{C})}{\lambda_0}, \quad \delta_1 = \frac{D_\kappa(C, \tilde{C})}{\lambda_1}, \quad (3.3)$$

for some fixed constants $\lambda_0, \lambda_1 > 0$, to be specified later.

This procedure may well decompose C and \tilde{C} into a different number of approximate bivertex arcs. If this occurs, it often indicates that the curves are not rigidly equivalent. However, this may also be due to noise, so we follow a procedure to delete less important arcs from the larger bivertex decomposition in this situation. Suppose the larger decomposition contains k more arcs than the smaller. Fixing some $l \geq 0$, the candidates for deletion will be from among the $k + l$ arcs on which curvature changes the least. For each of the $\binom{k+l}{k}$ possible choices of k arcs among these, we apply the comparison algorithm described below, and choose the one that yields the highest μ score. The value of l is used to control the extent to which we check less likely possibilities, and it is often reasonable to set $l = 0$, in which case the deletion is immediate. By following this procedure, we can henceforth restrict our attention to approximate bivertex decompositions of both curves that contain the same number of arcs.

3.2 Comparing Bivertex Arcs

The next task is to devise a method to compare individual bivertex arcs. Let B and \tilde{B} be two (approximate) bivertex arcs, with respective signatures $S(B)$ and $S(\tilde{B})$. We will construct a *signature similarity coefficient*, $p = p(B, \tilde{B}) \in [0, 1]$, to measure the closeness of the two signatures, where a p score of 1 reflects identical signatures and p decreases to 0 as the signatures become increasingly disparate.

Our motivating idea is to regard the two signature arcs as uniform wires that have opposite electrical charges, and then compute their mutual attraction, cf. [5, 14]. (Or, equivalently, their gravitational attraction assuming uniform density.) The choice of uniform (charge) density may well be modified, but so far this simplifying assumption has sufficed in our applications. The larger the attraction, the closer the signatures are to each other, and the more likely that the original arcs are congruent. However, rather than work directly with the signature arcs, we will employ their discretizations, denoted $S^\Delta, \tilde{S}^\Delta$, obtained either by sampling the full signatures, or by use of suitable numerical approximations to the values of the curvature invariants. We thus view the discretized signature as finite collections of oppositely charged points, and calculate their mutual attraction.

Let $\Sigma = \mathbb{R}^2$ indicate the *signature space* with coordinates κ, κ_s . Define the *separation* between two points $\sigma, \tilde{\sigma} \in \Sigma$ to be

$$d(\sigma, \tilde{\sigma}) = \begin{cases} \frac{\|\sigma - \tilde{\sigma}\|}{D - \|\sigma - \tilde{\sigma}\|}, & \|\sigma - \tilde{\sigma}\| < D, \\ \infty, & \|\sigma - \tilde{\sigma}\| \geq D, \end{cases} \quad (3.4)$$

where $\|\cdot\|$ is taken as the usual Euclidean norm³ on Σ , and where $D = D_\kappa(C, \tilde{C})$ is as defined in (3.2). Observe that $d(\sigma, \tilde{\sigma})$ is non-negative, and vanishes only when $\sigma = \tilde{\sigma}$, so that it defines a semimetric. Our analog of the electrical attraction between individual signature points $\sigma, \tilde{\sigma} \in \Sigma$ will be called the *strength of correspondence*, and inversely proportional to some power of their separation:

$$h(\sigma, \tilde{\sigma}) = \begin{cases} \frac{1}{d(\sigma, \tilde{\sigma})^\gamma + \epsilon}, & d(\sigma, \tilde{\sigma}) < \infty, \\ 0, & d(\sigma, \tilde{\sigma}) = \infty. \end{cases} \quad (3.5)$$

Here $\gamma > 0$ is fixed, and $\epsilon > 0$ is a small constant that serves as a cut-off to avoid infinities, so that $h(\sigma, \tilde{\sigma}) < \infty$ even when the signature points coincide.

Now, the strength of correspondence of two discretized signatures will be obtained by combining all the individual strengths $h(\sigma^i, \tilde{\sigma}^j)$ over all possible pairs of discrete signature points $\sigma^i \in S^\Delta$ and $\tilde{\sigma}^j \in \tilde{S}^\Delta$. We will need to appropriately rescale in order that the result lies in the interval $[0, 1]$. These will be based on the *rescaling function*

$$r(t) = \frac{t}{t + C_1}, \quad \text{for } 0 \leq t \leq \infty, \quad (3.6)$$

where $C_1 > 0$ is a constant, to be fixed later, that will influence the ultimate distribution of p scores along the interval $[0, 1]$. There are three evident ways to compute the p scores:

³As noted above, one could experiment with other (weighted) norms, but the Euclidean norm suffices for our purposes.

1. We rescale each $h(\sigma^i, \tilde{\sigma}^j)$ to $[0, 1]$ and let $p(S^\Delta, \tilde{S}^\Delta)$ be their average over all distinct pairs (i, j) .
2. We sum $h(\sigma^i, \tilde{\sigma}^j)$ over all pairs (i, j) and rescale the result to $[0, 1]$ to determine $p(S^\Delta, \tilde{S}^\Delta)$.
3. For each fixed i , we sum $h(\sigma^i, \tilde{\sigma}^j)$ over j and scale the result to $[0, 1]$. We then take $p(S^\Delta, \tilde{S}^\Delta)$ to be the average over i of the rescaled values.

The first two lead to comparison functions that are symmetric in S^Δ and \tilde{S}^Δ , while the third does not. We will see, however, that each of the symmetric methods harbors a not easily mended flaw, whereas the third, non-symmetric method, appears to avoid such pitfalls. Moreover, it can easily be made symmetric by a standard trick.

Consider procedure 1: If two bivertex arcs B and \tilde{B} are rigidly equivalent, then their signatures $S(B) = S(\tilde{B})$ are identical. However, there will exist pairs of discrete signature points $\sigma^i \in S^\Delta$ and $\tilde{\sigma}^j \in \tilde{S}^\Delta$ such that $d(\sigma^i, \tilde{\sigma}^j) > 0$. Therefore $p(S^\Delta, \tilde{S}^\Delta) < 1$, even though the curves are rigidly equivalent. This in itself is less than ideal, but more worrisome is the possibility that different pairs of identical curves might score differently. This causes us to abandon this approach.

Consider procedure 2: Suppose the discretized bivertex arc signatures $S^\Delta, \tilde{S}^\Delta$ contain a pair of almost identical signature points: $\sigma^i \approx \tilde{\sigma}^j$ for some i, j — which is not uncommon. Then, by this procedure, $h(\sigma^i, \tilde{\sigma}^j)$ will be very large, depending on the size of ε in (3.5), and thus $p(S^\Delta, \tilde{S}^\Delta) \approx 1$, regardless of the other data points.

Procedure 3 avoids these difficulties, and we now concentrate on developing it in more detail. Given $\sigma^i \in S^\Delta$, we set

$$h(\sigma^i, \tilde{S}^\Delta) = \sum_{j=1}^{\tilde{N}} h(\sigma^i, \tilde{\sigma}^j), \quad (3.7)$$

where $\tilde{N} = \#\tilde{S}$. Using (3.6) to rescale the result to $[0, 1]$ produces

$$p(\sigma^i, \tilde{S}^\Delta) = r(h(\sigma^i, \tilde{S}^\Delta)) = \frac{h(\sigma^i, \tilde{S}^\Delta)}{h(\sigma^i, \tilde{S}^\Delta) + C_1}. \quad (3.8)$$

At this point, we could compute $p(\sigma^i, \tilde{S}^\Delta)$ for each discrete signature point, and then average the results to produce an overall p score. However, we can do better. The Theorem of Turning Tangents, [6], says that for a simple closed curve C with positive orientation,

$$\oint_C \kappa(s) ds = 2\pi. \quad (3.9)$$

Thus, in a certain sense, points where $|\kappa(s)|$ is large determine a boundary curve more than do points with small curvature. Indeed, if we imagine modifying a curve by the addition of arcs, equation (3.9) stipulates that the insertion of an arc A requires the insertion of a second arc \tilde{A} with $\int_{\tilde{A}} \kappa(s) ds = -\int_A \kappa(s) ds$ in order that the resulting curve remain closed.

Arcs with zero total curvature, on the other hand, can be added with relative impunity. This motivates us to improve the comparison by weighting our final average by a power $\alpha > 0$ of the absolute value of the curvature at the discrete signature points $\sigma^i = (\kappa^i, \kappa_s^i) \in S^\Delta$, leading to

$$\widehat{p}(S^\Delta, \widetilde{S}^\Delta) = \frac{\sum_{i=1}^N p(\sigma^i, \widetilde{S}^\Delta) |\kappa^i|^\alpha}{\sum_{i=1}^N |\kappa^i|^\alpha}. \quad (3.10)$$

To impose symmetry on the final comparison function, we merely apply the method with the roles of the curves interchanged, and set

$$p(S^\Delta, \widetilde{S}^\Delta) = \min\{\widehat{p}(S^\Delta, \widetilde{S}^\Delta), \widehat{p}(\widetilde{S}^\Delta, S^\Delta)\}. \quad (3.11)$$

The use of the minimum is justified by the following argument. If C and \widetilde{C} are nearly rigidly equivalent, then $1 - \varepsilon \leq \widehat{p}(S^\Delta, \widetilde{S}^\Delta), \widehat{p}(\widetilde{S}^\Delta, S^\Delta) \leq 1$ for some small $\varepsilon > 0$, and so $0 \leq |\widehat{p}(S^\Delta, \widetilde{S}^\Delta) - \widehat{p}(\widetilde{S}^\Delta, S^\Delta)| \leq \varepsilon$. Thus, for nearly rigidly equivalent curves,

$$p(S^\Delta, \widetilde{S}^\Delta) = \min\{\widehat{p}(S^\Delta, \widetilde{S}^\Delta), \widehat{p}(\widetilde{S}^\Delta, S^\Delta)\} \approx \widehat{p}(S^\Delta, \widetilde{S}^\Delta) \approx \widehat{p}(\widetilde{S}^\Delta, S^\Delta).$$

On the other hand, if C and \widetilde{C} clearly differ, $|\widehat{p}(S^\Delta, \widetilde{S}^\Delta) - \widehat{p}(\widetilde{S}^\Delta, S^\Delta)|$ may be large, and (3.11) ensures that a pair of curves that mistakenly scored highly are reduced to a more appropriate p score.

3.3 Reconstructing Rigid Motions

Now, consider two general v -regular curves C and \widetilde{C} . As discussed above, we can assume that their approximate bivertex decompositions contain an equal number of bivertex arcs B_k, \widetilde{B}_k , for $k = 1, \dots, n$. We assume that the arcs are labelled sequentially and periodically as the curves are traversed with a counterclockwise orientation, so that $B_{k+n} = B_k$ and $\widetilde{B}_{k+n} = \widetilde{B}_k$ for all k . As a result, we need only compare B_k with \widetilde{B}_{k+m} for each $m = 0, \dots, n-1$. For each possible bijection, we take an average of the p scores of the corresponding bivertex arcs:

$$p_m(C, \widetilde{C}) = \frac{1}{n} \sum_{k=0}^{n-1} p(B_k, \widetilde{B}_{k+m}) = \frac{1}{n} \sum_{k=0}^{n-1} p(S_k^\Delta, \widetilde{S}_{k+m}^\Delta), \quad (3.12)$$

where $S_k^\Delta, \widetilde{S}_k^\Delta$ are the corresponding discretized signature arcs. If all of the resulting p_m scores are low, we return $\mu = 0$ and claim that C and \widetilde{C} are not rigidly equivalent. On the other hand, for each m such that $p_m(C, \widetilde{C}) > p_0$ for some fixed threshold $p_0 > 0$, we reconstruct the transformations $g_j \in \text{SE}(2)$ guaranteed by Theorem 8. If, for a particular m , all the resulting rigid motions are close to each other, we contend that C and \widetilde{C} are (nearly) rigidly equivalent.

To determine similarity between two rigid motions, we characterize each rigid motion $g_j \in \text{SE}(2)$ by the three parameters θ_j, a_j, b_j , so that

$$g_j \cdot z = \begin{pmatrix} \cos \theta_j & -\sin \theta_j \\ \sin \theta_j & \cos \theta_j \end{pmatrix} z + \begin{pmatrix} a_j \\ b_j \end{pmatrix}, \quad \text{for } z \in \mathbb{R}^2, \quad (3.13)$$

where the angle θ_j is measured modulo 2π . The group parameters are reconstructed by solving equations involving the weighted centers of mass and tangents of the arcs:

$$\begin{aligned} \begin{pmatrix} \cos \theta_j & -\sin \theta_j \\ \sin \theta_j & \cos \theta_j \end{pmatrix} \frac{\sum_j |\kappa_s^j|^\beta z^j}{\sum_j |\kappa_s^j|^\beta} + \begin{pmatrix} a_j \\ b_j \end{pmatrix} &= \frac{\sum_j |\tilde{\kappa}_s^{j+m}|^\beta z^{j+m}}{\sum_j |\tilde{\kappa}_s^{j+m}|^\beta}, \\ \theta_j = \arg \frac{\sum_j |\tilde{\kappa}_s^{j+m}|^\beta \tilde{z}_s^{j+m}}{\sum_j |\tilde{\kappa}_s^{j+m}|^\beta} - \arg \frac{\sum_j |\kappa_s^j|^\beta z_s^j}{\sum_j |\kappa_s^j|^\beta}. \end{aligned} \quad (3.14)$$

These quantities are then used to calculate the final μ score:

$$\mu = 1 - C_2 \left[\frac{\max_{j,k} R(\theta_j, \theta_k)}{\pi} + \frac{\max_{j,k} |a_j - a_k|}{D_x} + \frac{\max_{j,k} |b_j - b_k|}{D_y} \right], \quad (3.15)$$

where C_2 is a suitable constant, $R(\cdot, \cdot)$ is the smallest difference between angles when computed modulo 2π , while

$$\begin{aligned} D_x &= \max \left\{ \max_{z \in C} x - \min_{z \in C} x, \max_{z \in \tilde{C}} x - \min_{z \in \tilde{C}} x \right\}, \\ D_y &= \max \left\{ \max_{z \in C} y - \min_{z \in C} y, \max_{z \in \tilde{C}} y - \min_{z \in \tilde{C}} y \right\}. \end{aligned} \quad (3.16)$$

are characteristic distances based on the overall horizontal and vertical extents of the two curves.

4 An Application

The comparison method described in the preceding section was designed to test the rigid equivalence of plane curves, including those whose Euclidean signatures share a common trace. In this section, we describe the results of applying the algorithm to classify a database of mathematically constructed curves. More extensive applications and additional developments for practical digital images can be found in [7].

The curve database consisted of six classes, each corresponding to a particular member of a family of curves constructed in [9]. Specifically, the six generating curves C_r were constructed in MATHEMATICA, based on the code provided in [9, Example 1], and using (in the notation of that paper) the parameter values $r \in \{0, .5, 1, 1.5, 2, 2.5\}$. The six representative curves are plotted in Figure 1, while Figure 2 plots their Euclidean signatures' common trace.

Our comparison method easily distinguishes these six basic curves, meaning that each scores much higher when compared to itself than to the others. But it might still be claimed that two different discretizations of the same curve would also end up comparing poorly to each other. Therefore, to fully test the method, from each of the six initial curves, an additional 100 curves were generated by the addition of simulated noise. The method of noise simulation was designed to add in small random perturbations while maintaining overall smoothness. To this end, we first discretize the curve to produce $C^\Delta = \{z^1, \dots, z^N\}$.

We then randomly select $j = 1, \dots, t = 20$ sample points $z^{i_j} \in C^\Delta$, unit vectors u_j , and magnitudes $r_j \in [0, .005L]$, where L is the overall length of the curve. The sample points were then translated based on t independent Gaussian distributions, centered at each z^{i_j} :

$$z^k \mapsto z^k + \sum_{j=1}^t r_j e^{-2m(k,i_j)^2} u_j, \quad (4.17)$$

where

$$m(k, i) = \frac{1}{.05N} \min\{|k - i|, N - |k - i|\}. \quad (4.18)$$

Figure 3 displays a reference curve along with seven image curves obtained through application of this noise simulation. Although the curves appear nearly identical, the effect of the noise is evident in their disparate Euclidean signatures.

Each of the 600 noisy curves was compared to the six generating curves using the comparison method described in Section 3, with parameter values $\lambda_0 = 20$, $\lambda_1 = 10$, $l = 2$, $\gamma = 5$, $\epsilon = 10^{-4}$, $C_1 = 1000$, $\alpha = 2$, $p_0 = .7$, $\beta = 10$, and $C_2 = \frac{1}{3}$. A noisy curve was then deemed congruent to the generating curve that received the highest μ score, and considered to be correctly classified if it was sorted into the class containing its generator.

This scheme was run five times for five randomly generated databases of 600 curves, with each curve discretely represented by $N = 500$ sample points. The classifications were performed on Intel[®] Core[™] 2 Duo E8500 3.17GHz processors and completed in an average of 2,211 seconds, or approximately 37 minutes. The method classified the databases (3,000 curves in total) with an average success rate of 93.8%.

The software used to compute these examples and jigsaw puzzle assembly will soon be available on the second author's web page.

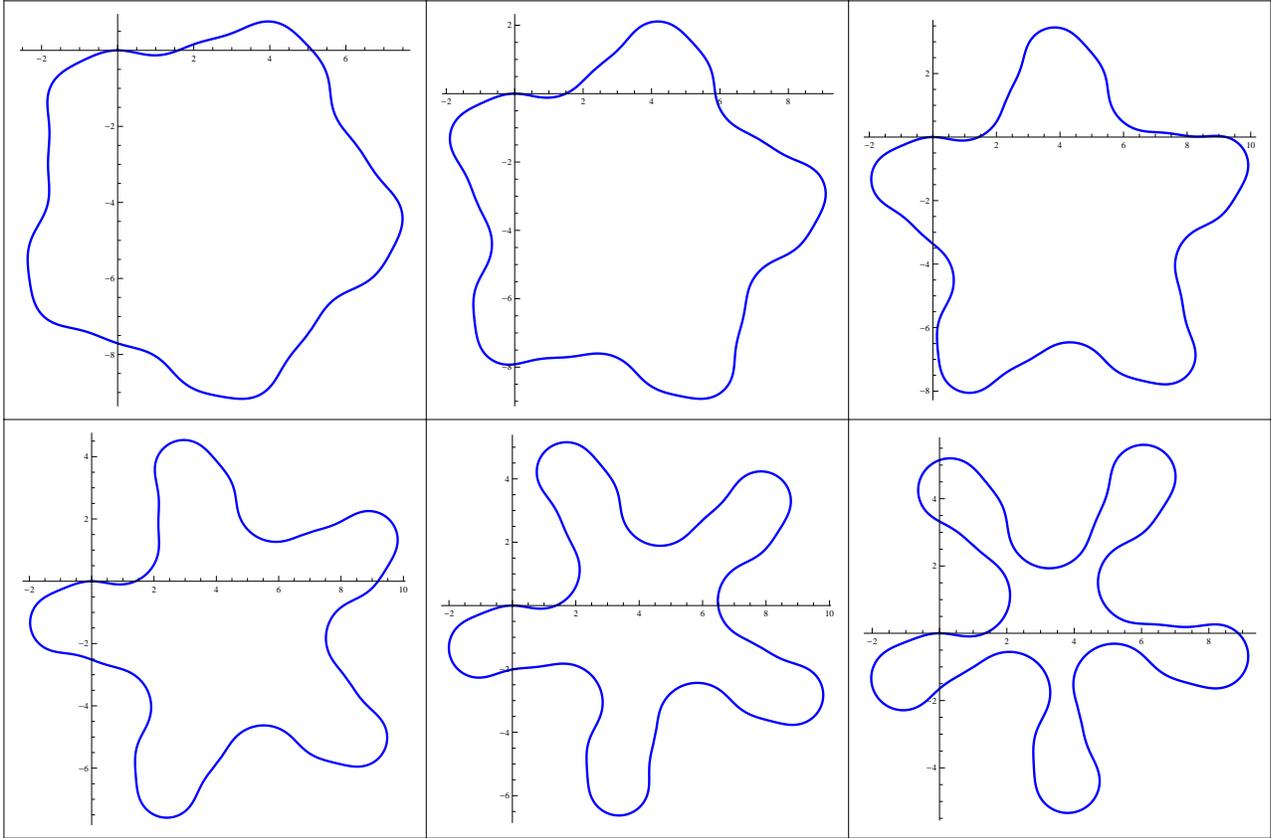


Figure 1: Plots of the six database generating curves, $\{C_r\}$. Curves were generated per the instructions provided in Example 1 from [9] for $r \in \{0, .5, 1, 1.5, 2, 2.5\}$.

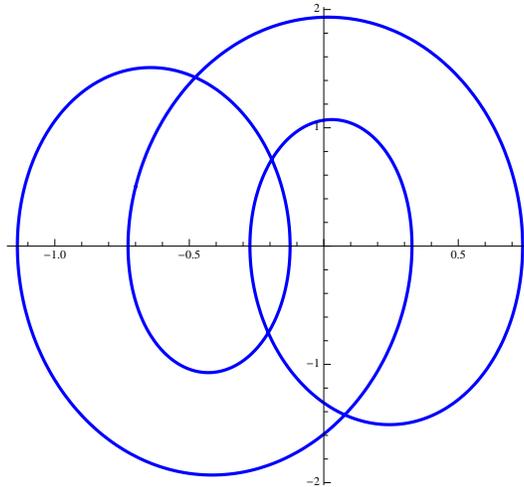


Figure 2: The common trace of the Euclidean signatures of the curves $\{C_r\}$ as developed in Example 1 from [9]. Six curves with signatures of this trace are plotted in figure 1.

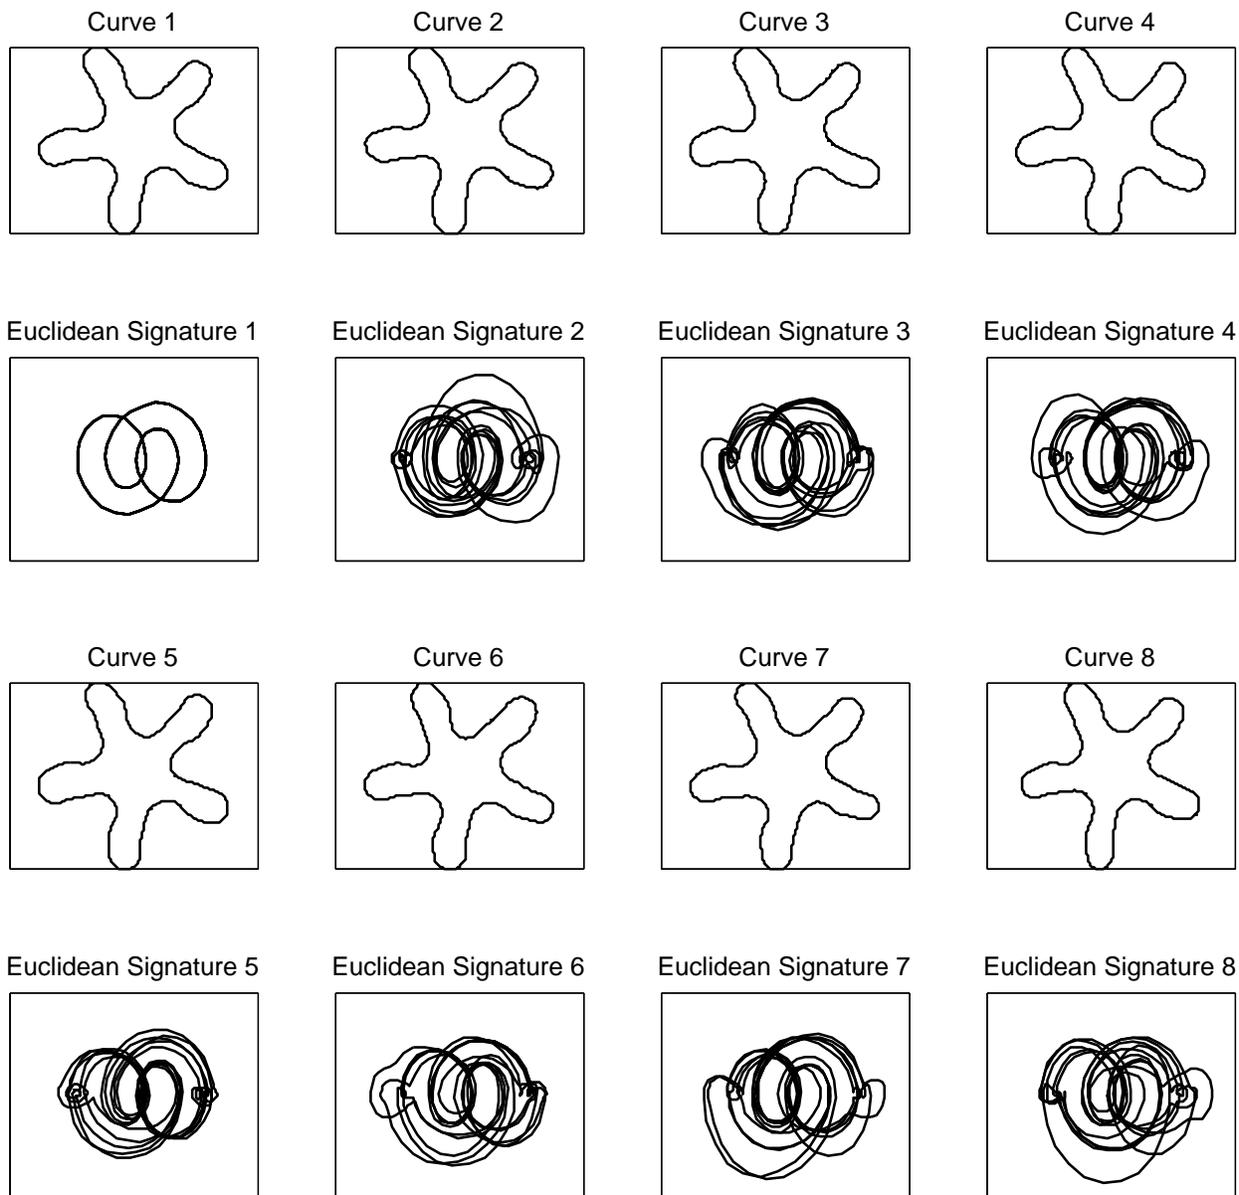


Figure 3: A demonstration of the method of noise simulation used to construct the database of curves. Curves 2-8 were developed by adding noise to curve 1. Notice that although the curves appear nearly identical, the noise is evident in their disparate Euclidean signatures.

References

- [1] Boutin, M., Numerically invariant signature curves, *Int. J. Computer Vision* **40** (2000), 235–248.
- [2] Calabi, E., Olver, P.J., Shakiban, C., Tannenbaum, A., and Haker, S., Differential and numerically invariant signature curves applied to object recognition, *Int. J. Computer Vision* **26** (1998), 107–135.
- [3] Cartan, É., Les problèmes d'équivalence, in: *Oeuvres Complètes*, part. II, vol. 2, Gauthier–Villars, Paris, 1953, pp. 1311–1334.
- [4] Fels, M., and Olver, P.J., Moving coframes. II. Regularization and theoretical foundations, *Acta Appl. Math.* **55** (1999), 127–208.
- [5] Freedman, M.H., Zheng-Xu He, Z.-X., and Wang, Z., Möbius energy of knots and unknots, *Ann. Math.* **139** (1994), 1–50.
- [6] Guggenheimer, H.W., *Differential Geometry*, McGraw–Hill, New York, 1963.
- [7] Hoff, D., and Olver, P.J., Automatic solution of jigsaw puzzles, preprint, University of Minnesota, 2011.
- [8] Hoffman, D.D., and Richards, W., Parts of recognition, *Cognition* **18** (1984), 65–96.
- [9] Musso, E., and Nicolodi, L., Invariant signature of closed planar curves, *J. Math. Imaging Vision* **35** (2009), 68–85.
- [10] Olver, P.J., Sapiro, G., and Tannenbaum, A., Differential invariant signatures and flows in computer vision: a symmetry group approach, in: *Geometry–Driven Diffusion in Computer Vision*, B.M. Ter Haar Romeny, ed., Kluwer Acad. Publ., Dordrecht, Netherlands, 1994, pp. 255–306.
- [11] Olver, P.J., *Equivalence, Invariants, and Symmetry*, Cambridge University Press, Cambridge, 1995.
- [12] Richards, W., Dawson, B., and Whittington, D., Encoding contour shape by curvature extrema, *J. Opt. Soc. Amer. A* **3** (1986), 1483–1491.
- [13] Wang, Y., *Smoothing Splines: Methods and Applications*, Monographs Stat. Appl. Probability, vol. 121, CRC Press, Boca Raton, FL, 2011.
- [14] Wu, K., and Levine, M.D., 3D part segmentation using simulated electrical charge distributions, *IEEE Trans. Pattern Anal. Machine Intel.* **19** (1997), 1223–1235.