

Integral Invariants for Shape Matching

Siddharth Manay, Daniel Cremers, *Member, IEEE*, Byung-Woo Hong, Anthony J. Yezzi Jr., *Member, IEEE*, and Stefano Soatto, *Member, IEEE*

Abstract—For shapes represented as closed planar contours, we introduce a class of functionals which are invariant with respect to the Euclidean group and which are obtained by performing integral operations. While such integral invariants enjoy some of the desirable properties of their differential counterparts, such as locality of computation (which allows matching under occlusions) and uniqueness of representation (asymptotically), they do not exhibit the noise sensitivity associated with differential quantities and, therefore, do not require presmoothing of the input shape. Our formulation allows the analysis of shapes at multiple scales. Based on integral invariants, we define a notion of distance between shapes. The proposed distance measure can be computed efficiently and allows warping the shape boundaries onto each other; its computation results in optimal point correspondence as an intermediate step. Numerical results on shape matching demonstrate that this framework can match shapes despite the deformation of subparts, missing parts and noise. As a quantitative analysis, we report matching scores for shape retrieval from a database.

Index Terms—Integral invariants, shape, shape matching, shape distance, shape retrieval.

1 INTRODUCTION

GEOMETRIC invariance is an important issue in computer vision that has received considerable attention in the past. The idea that one could compute functions of geometric primitives of the image that do not change under the various nuisances of image formation and viewing geometry was appealing; it held potential for application to recognition, correspondence, 3D reconstruction, and visualization. The discovery that there exist no generic viewpoint invariants was only a minor roadblock, as image deformations can be approximated with homographies; hence, the study of invariants to projective transformations and their subgroups (affine, similarity, Euclidean) flourished. Toward the end of the last decade, the decrease in popularity of research on geometric invariance was sanctioned mostly by two factors: the progress on multiple view geometry (one way to achieve viewpoint invariance is to estimate the viewing geometry) and noise. Ultimately, algorithms based on invariants did not meet expectations because most entailed computing various derivatives of measured functions of the image (hence, the name “differential invariants”). As soon as noise was present and affected the geometric primitives computed from the images, the invariants were dominated by the small-scale perturbations. Various palliative measures were taken, such as the introduction of scale-space smoothing, but a more

principled approach has so far been elusive. Nowadays, the field is instead engaged in searching for invariant (or insensitive) measures of photometric (rather than geometric) nuisances in the image formation process. Nevertheless, the idea of computing functions that are invariant with respect to group transformations of the image domain remains important because it holds the promise to extract compact, efficient representations for shape matching, indexing, and ultimately recognition.

1.1 Why Shape Distances?

Our ultimate goal is to compare objects represented as closed planar contours. This has obvious implications in shape classification for object recognition, content-based image retrieval, medical diagnosis, etc. At this level of generality, this is a monumental task that admits no simple meaningful solution [55]. Therefore, before we proceed any further, we need to specify what we mean by “objects,” how we describe their “shape” and concentrate our attention on particular ways in which they can “differ.” Within the scope of this paper, by *object* we mean a closed planar contour¹ with no self-intersections embedded in \mathbb{R}^2 ; its *shape* is the equivalence class of objects obtained under the action of a finite-dimensional group, such as the Euclidean, similarity, affine, or projective group [42]. In other words, two objects have the same shape if and only if one can be generated by rotating, translating and scaling the other. However, in comparing shapes, we want to be insensitive to certain variations that can occur to an object; for instance, in Fig. 1, we want rotated, jagged, articulated, and occluded hands to be judged as having shapes that are similar to that of the original hand. We prefer not to use the word “noise” when referring to these variations because, with the exception of the jaggedness, they are not obtained with standard additive, zero-mean, or small variance perturbations. For the case of the articulated and occluded objects, for instance, the perturbation can be quite significant in energy and highly localized along the contour. Our goal is to define a

- S. Manay is with the Electronics Engineering Technologies Division, Lawrence Livermore National Laboratory, PO Box 508, L-290 Livermore, CA 94551-0508. E-mail: smanay.ece98@gtalumni.org.
- D. Cremers is with the Department of Computer Science, University of Bonn, Roemerstrasse 164, D-53117 Germany. E-mail: dcremers@cs.uni-bonn.de.
- B.-W. Hong and S. Soatto are with the Computer Science Department, University of California Los Angeles, 3811 Boelter Hall, Los Angeles, CA 90095. E-mail: hong@cs.ucla.edu, soatto@ucla.edu.
- A.J. Yezzi Jr. is with the School of Electrical and Computer Engineering, Georgia Institute of Technology, 777 Atlantic Drive NW, Atlanta, GA 30332-0250. E-mail: ayezzi@ece.gatech.edu.

Manuscript received 12 Apr. 2005; revised 16 Feb. 2006; accepted 21 Feb. 2006; published online 11 Aug. 2006.

Recommended for acceptance by Y. Amit.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-0197-0405.

1. Many of our considerations can be extended to compact surfaces embedded in \mathbb{R}^3 .



Fig. 1. Variations of a sample shape.

shape distance so that shapes that vary by Euclidean transformations have zero distance and shapes that vary by scaling and moderate articulation and occlusion have small distance. In Fig. 1, all the hands should have low distance to each other, but high distance to other classes of shapes.

The type of variations we want to be resistant to can be lumped in three categories: “small deformations” that result in small set-symmetric differences between the interior of the curves being compared, “high-frequency noise” that affects a large portion of the contour and “localized changes” that significantly affect the total arclength of the contour but are spatially localized, such as spikes or wedges. Many researchers have addressed the comparison of objects under small deformations in a way that is invariant with respect to various transformation groups (see Section 2); fewer have addressed the sensitivity to high-frequency noise and yet fewer have addressed localized changes [76], [65]. In this paper, we plan to develop a framework that will address all of these variations simultaneously. To this end, we plan to employ a representation of shape in terms of *integral invariants*, so that the distance between objects will by construction be invariant with respect to the action of the chosen group; basing these invariants on integral computations allows us to address high-frequency noise in a principled way. Finally, establishing *point correspondence* (also called *shape matching*) among contours allows us to handle localized changes and small deformations. All these approaches are integrated into a shape distance that is insensitive to all these nuisances.

1.2 Differential versus Integral Invariants

Commonly, shape invariants are defined via differential operations. As a consequence, they are inherently sensitive to noise. As most practical applications of invariants require some robustness to small perturbations of the shape, it is necessary to revert to smoothing and accept the unfortunate side effect that meaningful information will be lost as well. In this work, we introduce invariants which are defined as integral functions of the shape. We restrict our analysis to Euclidean invariants, although extensions to the similarity and affine groups fit within the framework we propose. These integral invariants share the nice features of their differential counterparts, being invariant to certain group transformations and being local descriptors, which make them well-suited for matching under occlusions. Yet, in contrast to the differential invariants, the integral ones are inherently robust to noise and, therefore, do not require any preprocessing of the input data. In addition, they have the favorable feature that varying the size of the integration kernel provides a natural multiscale notion that, unlike differential scale spaces, does not require destructive smoothing.

1.3 From Invariants to Shape Distances

Based on integral invariants, we define a *shape distance* between matching parts. Here, a meaningful shape matching, a dense correspondence mapping the parametrized domains of one shape to another (and vice versa), is crucial, as distance

is defined as the integral (over the shape) of the difference between the invariant values of corresponding points. By minimizing an appropriate energy functional, we compute the optimal correspondence, which is affected both by differences in the local geometry of the two curves and by the amount of stretching or shrinking of the shapes’ parametrization required to map similar points to each other. Given this dense correspondence, the concepts of shape comparison, modeling, and interpolation can be naturally derived. We compute the optimal correspondence by casting the problem as one of identifying the shortest path in a graph structure, the nodes of which label possible correspondences between the points of the two contours. Similar shortest-path concepts were exploited in the context of shape matching and warping in [65], [76], [30], [41], [84], [4], [85].

In this paper, we briefly review the literature on shape analysis in this context (Section 2) before defining integral invariants and giving a few examples (Section 3). We then explore an optimization framework for computing shape distance and shape matching from invariants (Section 4), and detail the implementation of this framework (Section 5). In Section 4.1, we discuss the extension of the proposed integral invariants to multiscale analysis. Finally, before concluding, we demonstrate our method for computing correspondence and shape distance on noisy shapes (Section 6).

2 PREVIOUS WORK AND OUR CONTRIBUTION

Given the wealth of existing work on invariance, scale-space, and correspondence, our work naturally relates with a large body of literature, as we describe in the next subsection. The reader should notice that we consider each object as one entity and perform no analysis or decomposition, so there is no notion of hierarchy or compositionality in our representation, which is therefore intrinsically low-level.

2.1 Shape and Shape Matching

In the literature, one finds various definitions of the term shape. Kendall, for example, defines shape as whatever remains of point coordinates once you factor out a certain group transformation—for example, the similarity group covering translation, rotation, and scaling. We refer to [26] for a short review of the history of shape research. In this work, we refer to shape as a closed planar contour modulo certain group transformations. Moreover, we will denote by *shape matching* the process of putting into correspondence different parts of two given shapes. Applications of shape matching in computer vision include the classification of objects and the retrieval of objects of the same class based on the similarity of the object boundary [41]. In medical imaging, a given anatomical structure may be modeled by a statistical shape representation [45], [21]. Statistical representations of shape may also be useful when modeling complex shape deformations; for example, when observing the silhouette of a 3D object in various 2D views [18]. Intermediate shapes between two objects can be generally obtained based on their correspondence [76].

There exists a vast literature on comparing shapes, represented as a collection of points [3], [75], [87], [41], curves [48], [91], [64], [5], [30], [92], [77], and continuous curves reduced to various types of graph representations [94], [78], [80], [63], [49], [38]; we represent curves as continuous objects living in infinite-dimensional spaces. (In Section 5, we sample

the curve for implementation purposes.) Within this choice, many have addressed matching curves under various types of motions [3], [75], [87] and deformations [48], [91], [64], [5], [30], [22], [92], [15], [82], [77], some involving a mapping from one curve to another that has some optimality property [6], [30], [48], [64], [91], [15], [82], [5], [92], [85], [77].

The role of invariants in computer vision has been advocated for various applications ranging from shape representation [57], [7] to shape matching [6], [46], quality control [88], [14], and general object recognition [66], [1]. Consequently, a number of features that are invariant under specific transformations have been investigated [24], [39], [25], [33], [56], [81], [74].

In particular, one can construct primitive invariants of algebraic entities such as lines, conics, and polynomial curves, based on a global descriptor of shape [59], [28].

In addition to invariants to transformation groups, considerable attention has been devoted to invariants with respect to the geometric relationship between 3D objects and their 2D views. Invariant features can be computed from a collection of coplanar points or lines [67], [68], [32], [9], [29], [95], [1], [79], [40].

An invariant descriptor of a collection of points that relates to our approach is the shape context introduced by Belongie et al. [6], which consists of a radial histogram of the relative coordinates of the rest of the shape at each point.

Differential invariants to actions of various Lie groups have been addressed thoroughly [44], [36], [17], [58], [76], [30], [48], [64], [91], [85], [37]. An invariant is defined by an unchanged subset of the manifold which the group transformation is acting on. In particular, an invariant signature which pairs curvature and its first derivative avoids parameterization in terms of arc length [13], [60]. Calabi and coworkers suggested numerical expressions for curvature and first derivative of curvature in terms of joint invariants. However, it is shown that the expression for the first derivative of curvature is not convergent and modified formulas are presented in [8].

In order to reduce noise-induced fluctuations of the signature, semidifferential invariant methods are introduced by using first derivatives and one reference point instead of curvature, thus avoiding the computation of high-order derivatives [62], [31], [43]. Another semi-invariant is given by transforming the given coordinate system to a canonical one [89].

A useful property of differential and (some) semi-differential invariants is that they can be applied to match shapes despite occlusions, due to the locality of the signature [11], [10]. However, the fundamental problem of differential invariants is that high-order derivatives have to be computed, amplifying the effect of noise. There have been several approaches to decrease sensitivity to noise by employing scale-space via linear filtering [90]. The combination of invariant theory with geometric multiscale analysis is investigated by applying an invariant diffusion equation for curve evolution [71], [72], [16]. A scale parameter is another way to build a scale-space which is determined by the size of the differencing interval used to approximate derivatives using finite differences [12]. In [54], a curvature scale-space was developed for a shape matching problem. A set of Gaussian kernels was applied to build a scale-space of curvature whose extrema were observed across scales.

To overcome the limitations of differential invariants, there have been attempts to derive invariants based on integral computations. A statistical approach to describe invariants was introduced using moments in [35]. Moment invariants under affine transformations were derived from the classical moment invariants in [27]. They have a limitation in that high-order moments are sensitive to noise which results in high variances. The error analysis and analytic characterization of moment descriptors were studied in [47]. The Fourier transform was also applied to obtain integral invariants [93], [52], [2]. A closed curve was represented by a set of Fourier coefficients and normalized Fourier descriptors were used to compute affine invariants. In this method, high-order Fourier coefficients are involved and they are not stable with respect to noise. Several techniques have been developed to restrict the computation to local neighborhoods. The wavelet transform was used for affine invariants using the dyadic wavelet in [83] and potentials were also proposed to preserve locality [34]. Alternatively, semilocal integral invariants are presented by integrating object curves with respect to arc length [73]. More recently, attempts to develop invariants with the locality properties, but without the sensitivity, of differential invariants have resulted in functions of curves that are based not on differential operators, but on integral operators applied to the contour or the characteristic function of its interior [50], [65].

In this paper, we introduce a framework for integral invariants along with two general classes of integral invariants. We use the resulting invariant descriptions to define a notion of distance between shapes and we illustrate the potential of our representation on several experiments with simulated and real images.

2.2 Implicit versus Explicit Contour Representations

In the context of image segmentation, the *implicit* representation of closed contours as the zero-crossing of corresponding embedding functions has become increasingly popular. The level set method [23], [61] provides a framework to elegantly propagate boundaries in a way which allows for topological changes of the embedded contour and does not require reparameterization. Recently, shape dissimilarity measures and statistical shape models have been formulated on the basis of the level set representation [45], [86], [69], [20], [19]. Yet, such implicit representations do not provide inherent support for pointwise correspondences. In order to model the notion of corresponding features and parts and, therefore, take these correspondences into account in a model of shape similarity (quantified by shape distance), we revert to *explicit* parameterizations of closed contours.

3 INTEGRAL INVARIANTS

In this section, we focus on the definition and examples of integral invariants.

Throughout this section, we indicate with $C : \mathbb{S}^1 \rightarrow \mathbb{R}^2$ a closed planar contour with infinitesimal arclength ds and G a group acting on \mathbb{R}^2 , with dx the area form on \mathbb{R}^2 . We also use the formal notation \bar{C} to indicate the interior of the region bounded by C . μ is either the curve C itself (a one-dimensional object) or \bar{C} (a two-dimensional object) and

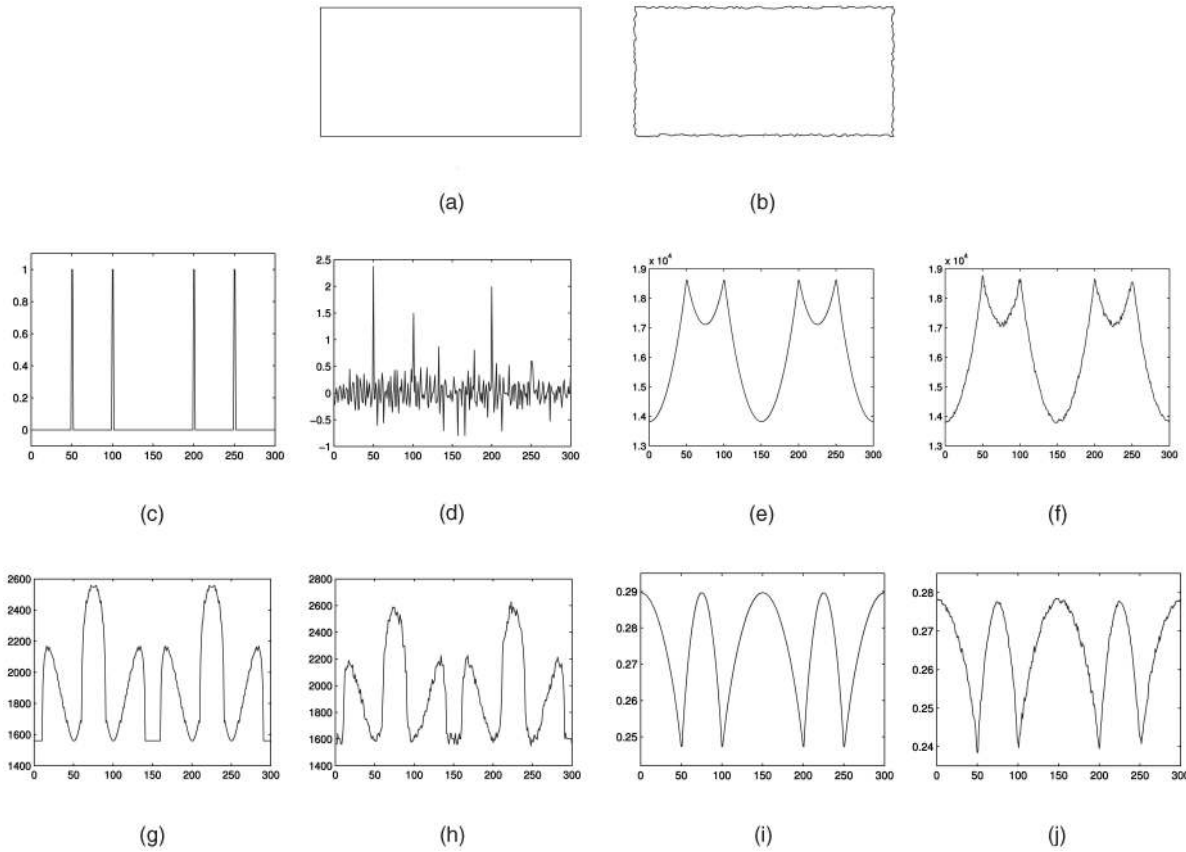


Fig. 2. Demonstration of the effect of noise on different invariants. (a) A rectangular shape. (b) A rectangular shape with noise. (c) Curvature of (a). (d) Curvature of (b). (e) Distance Integral Invariant of (a). (f) Distance Integral Invariant of (b). (g) Local Distance Integral Invariant of (a). (h) Local Distance Integral Invariant of (b). (i) Local Area Integral Invariant of (a). (j) Local Area Integral Invariant of (b).

$d\mu(x)$ the corresponding measure, i.e., the area form dx or the infinitesimal arclength ds , respectively.

Definition 1. Let G be a transformation group acting on \mathbb{R}^2 . A function $I : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a G -invariant if it satisfies

$$I(C) = I(g \cdot C), \quad \forall g \in G. \quad (1)$$

The function $I(\cdot)$ associates to each point on the contour a real number. In particular, if the point $p \in C$ is parameterized by arclength, the invariant can be interpreted as a function from $[0, L]$, where L is the length of the curve, to the reals:

$$\{C : \mathbb{S}^1 \rightarrow \mathbb{R}^2\} \mapsto \{I_C(p(s)) : [0, L] \rightarrow \mathbb{R}\}. \quad (2)$$

Similarly, if $p \in C$ is parameterized from $[0, 1]$, the invariant can be interpreted as a function from $[0, 1]$ to the reals:

$$\{C : \mathbb{S}^1 \rightarrow \mathbb{R}^2\} \mapsto \{I_C(p(s)) : [0, 1] \rightarrow \mathbb{R}\}. \quad (3)$$

We abuse this generalized notation in our discussions.

This formal definition of an invariant includes some very familiar examples, such as curvature.

Example 1 (Curvature). For $G = SE(2)$, the curvature κ of C is G -invariant.

The profiles of the curvature for the rectangular shape in Fig. 2a and its noisy version in Fig. 2b are shown in Figs. 2c and 2d, respectively. The curvature is called *differential invariant* since its calculation is based on differential operations. The curvature is a useful feature for describing shapes at matching due to its invariant property under a group

transformation of $SE(2)$, which will be considered as a transformation group G for the following invariants. However, as shown in Fig. 2d, it is sensitive to noise because the calculation of the curvature is dependent on second-order derivatives. Thus, we introduce an invariant that is robust to noise by employing integral operations for its calculation. We begin with a general notion of integral invariant.

Definition 2. A function $I_C(p) : \mathbb{R}^2 \rightarrow \mathbb{R}$ is an integral G -invariant if there exists a kernel $k : \mathbb{R}^2 \times \mathbb{R}^2 \rightarrow \mathbb{R}$ such that

$$I_C(p) = \int_{\mu} k(p, x) d\mu(x), \quad (4)$$

where $k(\cdot, \cdot)$ satisfies

$$\int_{\mu} k(p, x) d\mu(x) = \int_{g\mu} k(gp, x) d\mu(x) \quad \forall g \in G, \quad (5)$$

$$g\mu \doteq \{gx \mid g \in G, x \in \mu\}.$$

The definition can be extended to vector invariants² or to multiple integrals. Note that the point p does not necessarily lie on the contour C , as long as there is an unequivocal way of associating $p \in \mathbb{R}^2$ to C (e.g., the centroid of the curve).

Note that a regularized version of curvature, or, in general, a curvature scale-space, can be interpreted as an integral

2. While not discussed in this paper, vector-valued invariants could be composed of, e.g., integral invariants at multiple scales, invariants of increasing derivative or integral order, or even unrelated invariants.

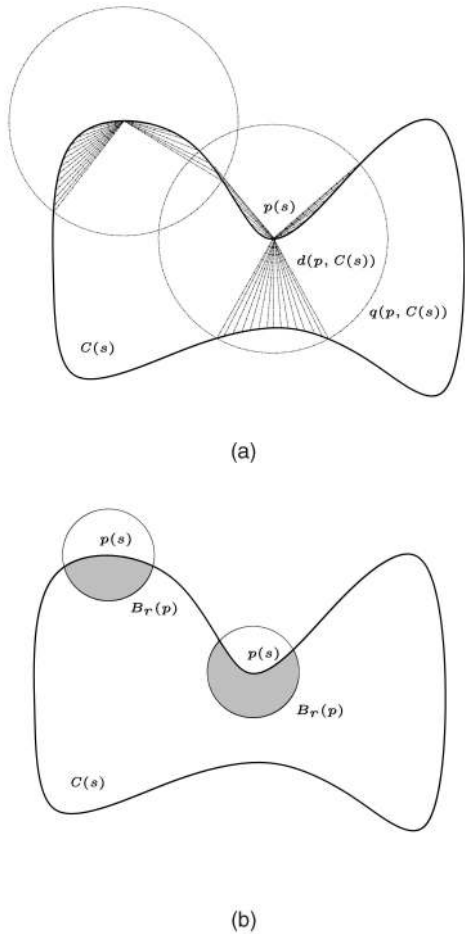


Fig. 3. (a) Distance integral invariant defined in (6), made local by means of a kernel as described in (7). (b) Integral local area invariant defined by (9).

invariant since regularized curvature is an algebraic function of the first and second-regularized derivatives [54]. Therefore, integral invariants are quite general and contain regularized differential invariants as a subset. However, the spirit of this manuscript is to avoid the computation of derivatives of contour data, so we will not explore this subset.

Example 2 (Distance integral invariant). Consider $G = SE(2)$ and the following function, computed at every point $p \in C$:

$$I_C(p) \doteq \int_C |p - x| ds(x), \quad (6)$$

where $|y - x|$ is the Euclidean distance in \mathbb{R}^2 . This is illustrated in Fig. 3a.

It is immediate to show that this is an integral Euclidean invariant, since Euclidean transformations preserve distance. We note that, unlike curvature, the range of values for the distance invariant is \mathbb{R}^+ (since the Euclidean distance is always nonnegative).

The profiles of the distance integral invariant for the shapes in Figs. 2a and 2b are shown in Figs. 2e and 2f, respectively. The distance integral invariant is robust to noise, the effect of which is reduced, as shown in Fig. 2f. However, it is a global descriptor in that a local change of a shape affects the values of the distance integral invariant for the entire shape.




Fig. 4. (a) A rectangular shape with two mark points \circ and Δ . (b) Local Distance Integral Invariant of (a) and corresponding mark points, which have the same invariant value even though they have very different shapes (i.e., a corner and a straight line). (c) Local Area Integral Invariant of (a) and corresponding mark points.

A version of the invariant I_C that preserves locality can be obtained by weighting the integral in (4) with a kernel $q(p, x)$, so that $I_C(p) \doteq \int_C k(p, x) ds(x)$ where

$$k(p, x) \doteq q(p, x) d(p, x). \quad (7)$$

The kernel $q(\cdot, \cdot)$ is free for the designer to choose depending on the final goal. This local integral invariant can be thought of as a continuous generalization of “shape context,” which was designed for discretized shapes represented as points [6]. The difference is that the shape context signature is a local radial histogram of neighboring points, whereas in our case, we only store the mean of their distance. This allows extension to continuous representations of shapes and obviates the need to choose histogram granularity, etc.

The local distance integral invariant is a local descriptor provided by the integral kernel restricted on a circular neighborhood. It is also robust to noise as shown in Figs. 2g and 2h. Thus, it may be effective for both noise and occlusion. However, this invariant is not discriminative in that it can have the same value for different geometric features. This drawback is demonstrated in Fig. 4. The two points marked by \circ and Δ on different geometric features of the shape in Fig. 4a have the same local distance integral invariant as shown in Fig. 4b. This is a motivation to introduce the following invariant.

Example 3 (Area integral invariant). Define a ball $B_r(p)$ as a function $B_r : \mathbb{R}^2 \times \mathbb{R}^2 \rightarrow \{0, 1\}$ to be an indicator function on the interior of a circle with radius r centered at p ,

$$B_r(p, x) = \begin{cases} 1 & |p - x| \leq r \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

For any given radius r , the corresponding integral invariant,




Fig. 5. Bottom row: Multiscale local area integral invariant for the shapes (top row). (a) A rectangular shape. (b) A rectangular shape with noise.

$$I_C^r(p) \doteq \int_C B_r(p, x) dx, \tag{9}$$

can be thought of as a function from the interval $[0, L]$ to \mathbb{R}^+ (since area is always nonnegative), bounded above by the area of the region bounded by the curve C . This is illustrated in Fig. 3b and examples are shown in Fig. 2.

As shown in Figs. 2i and 2j, the local area integral invariant is robust to noise and has a locality property similar to the local distance integral invariant. In addition, it has a strong descriptive power with respect to the marked points on the shape shown in Fig. 4c. Thus, the local area integral invariant is an effective descriptor for shape matching and we rely on a variant of this integral invariant throughout this work.

We note again that this integral invariant is a case of “shape context,” this time with one histogram bin. In addition to allowing extension to continuous representations of shape, local area invariants are explicitly rotationally invariant, while shape context achieves rotational invariance by computing circular shifts of the radial descriptor.

Naturally, if we plot the value of $I_C^r(p(s))$ for all values of s and r ranging from zero to a maximum radius so that the local kernel encloses the entire curve $B_r(p) \supset C \quad \forall p \in C$ (at which point the invariant would be a constant), we can generate a graph of a function that can be interpreted as a multiscale integral invariant, as shown in Fig. 5. We will return to this idea in Section 4.1. Furthermore, $B_r(p, x)$ can be substituted by a more general kernel, for instance, a Gaussian centered at p with $\sigma = r$.

Note also that the integral invariant can also be defined as normalized by the area of $B_r(p)$ for convenience.

Example 4 (Normalized area integral invariant).

$$I_C^r(p) \doteq \frac{\int_C B_r(p, x) dx}{\int_{\mathbb{R}^2} B_r(p, x) dx}. \tag{10}$$

The corresponding integral invariant is then bounded between 0 and 1. Because it mimics the qualities of the Area Integral Invariant discussed above, this is the invariant we favor in the remainder of this work.




Fig. 6. Construction of approximation of the area integral invariant. (a) The computation of θ . (b) the approximation of $I_C^r(p)$.

3.1 Relation of Local Area Integral Invariant to Curvature

In this section, we note a formal connection between traditional differential invariants and the Local Area Integral Invariant in a limiting case. Curvature provides a useful descriptor for shape matching due to its invariance and locality. It is considered a *complete* invariant in the sense that it allows the recovery of the original curve up to the action of the symmetry group. Furthermore, all differential invariants of any order on the plane are functions of curvature [89] and, therefore, linking our integral invariant to curvature would allow us to tap into the rich body of results on differential invariants without suffering from the shortcomings of high-order derivatives at its computation.

We first assume that the curve C is smooth,³ so that a notion of curvature is well-defined and the curvature can be approximated locally by the osculating circle shown in Fig. 6. The invariant $I_r(p)$ denotes the area of the intersection of a circle with radius r with the interior of C , and it can be approximated to first-order by the area of the shaded sector in Fig. 6, i.e., $I_r(p) \simeq 2r^2\theta$. Now, the angle θ can be computed as a function of r and R using the cosine law: $\cos(\theta) = r/2R$. Since curvature κ is the inverse of R we have

$$I_r(p) \simeq 2r^2 \cos^{-1}\left(\frac{1}{2}r\kappa(p)\right). \tag{11}$$

Now, since $\cos^{-1}(x)$ is an invertible function, to the extent in which the approximation above is valid (which depends on r), we can recover curvature from the integral invariant. The approximation above is valid in the limit when $r \rightarrow 0$.

4 SHAPE MATCHING AND DISTANCE

Given two shapes represented by curves C_1, C_2 , we want to compute their shape distance, a scalar that quantifies the similarity of the two contours. Basing this computation on a group invariant will ensure that the shape distance is not affected by group actions on the shape; further basing it on an integral invariant will make the distance computation robust to noise and local deformations of the contour. Naively, we

3. Notice that our invariant does *not* require that the shape be smooth, and that this assumption is made only to relate our results to the literature on differential invariants.

could define the shape distance to be the difference between the invariant functions, but upon further reflection we see that this distance is meaningful only if the computation somehow compares similar parts of the two shapes. If we compare one (for example) rabbit's ears to another's leg, we will decide (incorrectly) that the two shapes are very different. Yet this will be the effect of computing a shape distance without first establishing a dense correspondence between the points of the contours. Computing the difference of invariant values between corresponding points implies comparing one rabbit's ears to the second rabbit's ears, a much more meaningful metric.

Thus, we wish to find an optimal correspondence between the contours and concurrently measure the shape distance based on the correspondence. We will express the correspondence with a continuous *disparity function* $d(s) : \mathbb{S} \mapsto \mathbb{R}$ that reparameterizes two curves C_1 and C_2 (and their integral invariants I_1 and I_2). Given two reparameterizations $h_1, h_2 : \mathbb{S} \mapsto \mathbb{S}$ where $C_1(h_1(s)) \sim C_2(h_2(s))$ (\sim denotes correspondence), the disparity function is

$$d(s) = \frac{h_1(s) - h_2(s)}{2}. \quad (12)$$

We require that h_1, h_2 are nondecreasing functions of s , from which we derive $-1 \leq d'(s) \leq 1$. Below, we define an energy functional $E(\dots, d)$; the disparity function d^* that minimizes an energy functional

$$d^*(s) = \arg \min_{d(s)} E(\dots, d) \quad (13)$$

describes the optimal point correspondence between the two curves,

$$C_1(s - d^*(s)) \sim C_2(s + d^*(s)), \quad \forall s \in \mathbb{S}. \quad (14)$$

Intuitively, two corresponding points on two contours should have similar invariant values, which leads us to define the energy functional $E(I_1, I_2, d)$ for the discrepancy between two integral invariants I_1, I_2 , in terms of the disparity function $d(s)$, as follows:

$$\begin{aligned} E(I_1, I_2, d) &= E_1(I_1, I_2, d) + \alpha E_2(d') \\ &= \int_0^1 \|I_1(s - d(s)) - I_2(s + d(s))\|^2 ds \\ &\quad + \alpha \|d'(s)\|^2 ds, \end{aligned} \quad (15)$$

where $\alpha > 0$ is a constant. The first term E_1 of the energy functional measures the similarity of two curves by integrating the local difference of the integral invariant at corresponding points. A cost functional based on a local comparisons minimizes the impact of articulations and local changes of a contour because the difference in invariants is proportionally localized in the domain of the integral; contrast this with a global descriptor where local changes influence the descriptor everywhere. Representing the correspondence with $d(s)$ ensures that the E_1 term is both symmetric (we make this precise below) and independent of the arclength of $d(s)$.⁴

4. Correspondence representations such as the "warping function" introduced in Section 5 implicitly couple the E_1 term to the correspondence's arclength and, thus, require normalization terms. We favor separating the energy functional into an arclength-free penalty computed on corresponding points in the E_1 term, and an explicit warping penalty in the E_2 term.




Fig. 7. Correspondences between two invariant signals $I_1(s)$ (top) and $I_2(s)$ (bottom) with different values for the control parameter α in the energy functional. Smaller values of the parameter α in (15) will facilitate contour shrinking and stretching in the matching process. (a) $\alpha = 100$, (b) $\alpha = 30$, and (c) $\alpha = 1$.

The second term E_2 of the energy functional is the elastic energy of the disparity function $d(s)$ that penalizes stretching or shrinking of the mapped curve length. When $d(s) = 0$, the parameterizations of the two contours instruct the matching directly (i.e., points on the contour with the same parameter value correspond). $d(s)$ such that $d'(s) = 0$ indicates circular "shifts" of the correspondence. Other values of $d(s)$ "stretch" or "shrink" the length of segments of one contour onto the other; it is this action of $d(s)$ that the E_2 energy term penalizes.

To demonstrate the effect of the control parameter α in the energy functional, one example of the optimal correspondence between two integral invariants with various values of the control parameter is shown in Fig. 7. One integral invariant is represented by a straight line shown on the bottom and the other integral invariant is represented by a line with a spike shown on the top in each figure. The larger the control parameter α , the more correspondence is regularized, as shown in Fig. 7a. Fig. 7c shows that a feature characterized by the spiculation in one integral invariant on the top is mapped to an infinitesimal portion in the other integral invariant on the bottom. The difference of geometrical features is emphasized more with a small α .

Ultimately, a notion of shape distance should be symmetric. It is generally undesirable to privilege one shape rather than the other when matching two shapes. The energy functional defined in (15) is designed to satisfy a symmetry property that gives

$$\begin{aligned} d^*(s) = \arg \min_{d(s)} E(I_1, I_2, d) &\iff d^*(s) = \arg \min_{d(s)} E(I_2, I_1, d) \\ E(I_1, I_2, d^*) &= E(I_2, I_1, -d^*). \end{aligned} \quad (16)$$

The shape distance $D(C_1, C_2)$ between two curves C_1, C_2 is measured via the optimal correspondence $d^*(s)$ in the energy functional E between their integral invariants I_1, I_2 as defined by

$$D(C_1, C_2) = E(I_1, I_2, d^*). \quad (17)$$




Fig. 8. Scalogram of the shape in Fig. 2a and trace of local extrema across scales. (a) Scalogram of the curvature scale-space. Note the dislocation of the extrema. (b) Scalogram of the integral invariant scale-space.

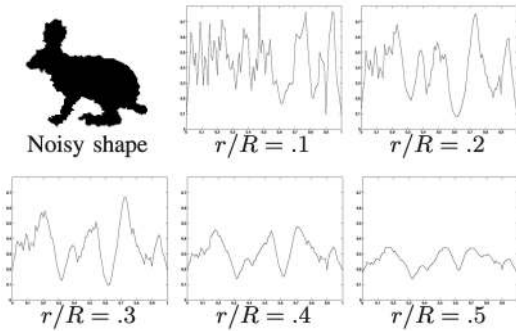


Fig. 9. Integral Invariant for a noisy shape computed at different scales. $R = 120$ for this shape.

Since the energy functional is symmetric (up to the sign of d^*), the shape distance is as well. In Section 5, we outline the computation of d^* and $D(C_1, C_2)$.

4.1 Shape Matching with Multiscale Integral Invariants

The integral invariant intrinsically introduces the notion of scale, varying the size of the kernel naturally forms a multiscale invariant. While we do not exploit the multiscale properties of these invariants or multiresolution methods in this work, this section highlights some of the features and potential of the integral invariant framework.

In [53], a curvature scale-space is used in a coarse-to-fine procedure to establish a sparse matching between inflection points of two shapes. Matches at a given scale are computed from matches at coarser scales. However, mismatching in the first stage causes fatal errors. Further, since curvature scale-space is derived from Gaussian smoothing, the inflection points move with increased blurring, and reparameterization is required to find correspondence between these inflection points across scale. In Fig. 8, the curvature scale-space and a multiscale integral invariant (more specifically, the local area invariant with varying kernel radius) for the shape in Fig. 2a are compared. The dislocation of the extrema points occurs across scales in the curvature scale-space, as shown in Fig. 8a. In contrast, the location of the extrema points stays the same across scales in the multiscale integral invariant, as in Fig. 8b since features at various scales are observed based on the original shape rather than on blurred




Fig. 10. Demonstration of correspondences between two rectangular shapes with spikes at different scales. The figures show the optimal point correspondence determined by our algorithm for increasing size of the kernel width r in (10). The two spikes are identified as “corresponding” on a fine scale only. (a) Fine scale, (b) intermediate scale, and (c) coarse scale.

versions of the shape. This obviates the need to compute correspondences between scales.

A multiscale integral invariant, demonstrated in Fig. 9, could also be used in a hierarchical description of features. The matching using the integral invariants at a fine scale provides a correspondence taking into account detail features on the shapes while matching at a coarse scale considers large features on the shapes. This is demonstrated in Fig. 10 on the fine scale peak, which is accounted for in the fine-scale matching but ignored in the coarser scale matchings.

With any method that allows a choice of scale, the question of which scale is optimal for shape analysis arises. While the choice of an optimal scale is a subject of continuing investigation, we offer the following comment. In Section 3, we noted that the kernel radius of the Local Area Integral Invariant is bounded above by a value R , for which the kernel encompasses the entire shape. At this scale, the value of the invariant is constant. So, rather than express scale in absolute terms parameterized by the kernel radius r , scale can also be parameterized in relative terms measured by the normalized kernel radius r/R , as in Fig. 9. This provides a common context to parameterize scale even among shapes of different sizes.

5 IMPLEMENTATION

In Section 4, we presented a distance between invariants (and therefore shape) that depends on the choice of a disparity function $d(s)$. To complete the calculation of distance and to establish a local correspondence between the curves, we must optimize distance with respect to $d(s)$. This section briefly outlines the implementation of the computation of the local area integral invariant and a well-known approach to globally optimize the correspondence for a discrete representation of the curves as ordered sets of points.

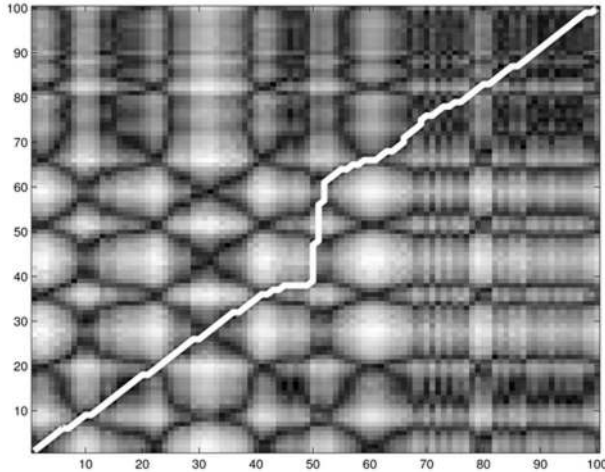


Fig. 11. Optimal path through the graph. The path warps the parameterization of the hand (on the bottom of the graph) and the parameterization of the noisy occluded (four-fingered) hand (on the left side of the graph). Both shapes are shown in Fig. 1. The gray levels indicate the dissimilarity between points; lighter shade indicates higher dissimilarity. See text for more details.

To efficiently compute the local area integral invariant, consider the binary image $\chi(\bar{C})$ (an indicator function defined as 1 on the interior of the contour and 0 elsewhere) and convolve it with the kernel $k(p, x) \doteq B_r(p - x)$, where $p \in \mathbb{R}^2$, not just the curve C . Evaluating the result of this convolution on $p \in C$ yields I_C , without the need to parameterize the curve. However, we retain a parametrized representation of the curve for the computation of the correspondence.

Our implementation is based on dynamic programming approaches similar to those employed by many in the shape, stereo, and registration (for medical imaging) communities [70], [65], [76], [30], [41], [84], [4], [85].

In order to adopt a graph search framework, the representation of the correspondence needs to be recast as a parameterized path. Thus, we exchange the disparity function $d(s)$ for a *warping function*,⁵ $h(\xi) = (h_1(\xi), h_2(\xi))$ for curves $(C_1(s), C_2(s))$ in the energy functional defined in (15) by setting

$$\begin{aligned} s - d(s) &= h_1(\xi) \\ s + d(s) &= h_2(\xi). \end{aligned} \quad (18)$$

Since $h_i : [0, L_h] \mapsto \mathbb{S}^1$ for $i \in \{1, 2\}$ (with $h_i(0) = h_i(L_h)$), the warping function $h : [0, L_h] \mapsto \mathbb{S}^1 \times \mathbb{S}^1$ is a path in $\mathbb{S}^1 \times \mathbb{S}^1$, parameterized by its length L_h , that represents the pointwise correspondence between the curves $C_1(s), C_2(s)$:

$$C_1(h_1(\xi)) \sim C_2(h_2(\xi)), \quad \forall \xi \in [0, L_h] \subset \mathbb{R}. \quad (19)$$

5. The representation of the reparameterization of two curves to achieve pointwise correspondence has been approached in several different ways. For example, in [76], symmetric and asymmetric “alignment functions” are used. All of these methods are somewhat similar, with one main distinction: symmetric representations result in symmetric energy and distance functions, while asymmetric representations do not. Due to our focus on shape distance, we restrict ourselves to symmetric reparameterizations: the disparity function because it results in an energy functional that does not depend implicitly on pathlength, and the warping function for its simplicity to implement in a discrete dynamic programming framework.




Fig. 12. Schematic view of a graph used to compute the correspondence for two curves with $M = N = 5$. The dots are the nodes, and the arrows are the directed edges. The dashed arrows are the “wrap around” edges. A typical portion of this graph is shown enlarged in Fig. 13.

Differentiating (18) with respect to s yields

$$\begin{aligned} 1 - d'(s) &= h'_1(\xi) \frac{d\xi}{ds} \\ 1 + d'(s) &= h'_2(\xi) \frac{d\xi}{ds}. \end{aligned} \quad (20)$$

Then, the original energy functional in (15) becomes

$$\begin{aligned} \tilde{E}(I_1, I_2, h_1, h_2) &= \\ &\int_0^{L_h} \left(\|I_1(h_1(\xi)) - I_2(h_2(\xi))\|^2 \right) \left(\frac{h'_1(\xi) + h'_2(\xi)}{2} \right) d\xi \\ &+ \alpha \int_0^{L_h} \left(\left\| \frac{h'_2(\xi) - h'_1(\xi)}{h'_1(\xi) + h'_2(\xi)} \right\|^2 \right) \left(\frac{h'_1(\xi) + h'_2(\xi)}{2} \right) d\xi. \end{aligned} \quad (21)$$

In this way, the warping function $h(\xi)$ derives a formula for the energy functional in terms of $(h_1(\xi), h_2(\xi))$,

$$h : E(I_1, I_2, d) \rightarrow \tilde{E}(I_1, I_2, h_1, h_2). \quad (22)$$

Then, finding an optimal disparity function $d^*(s)$ in the energy functional E becomes equivalent to finding an optimal warping function $h^*(\xi) = (h_1^*(\xi), h_2^*(\xi))$ in the energy functional \tilde{E} as follows:

$$\begin{aligned} d^*(s) &= \arg \min_{d(s)} E(I_1, I_2, d) \\ &\Updownarrow \\ (h_1^*(\xi), h_2^*(\xi)) &= \arg \min_{h_1(\xi), h_2(\xi)} \tilde{E}(I_1, I_2, h_1, h_2). \end{aligned} \quad (23)$$

To exploit the dynamic programming framework, we must discretize the curve by sampling it at uniform intervals. The result is an ordered set of points.

In the discrete case, an intuitive algorithm to compute shape matching would be as follows: We first find an initial correspondence between a point on each curve, say $C_1[i] \sim C_2[j]$ (more on this below). The “next” correspondence should be the choice of action that minimizes the energy ((15)); the possible actions are 1) locally contracting the first curve onto the second, setting $C_1[i + 1] \sim C_2[j]$, 2) locally contracting the second curve onto the first, setting $C_1[i] \sim C_2[j + 1]$, or 3) locally mapping the discrete points as one-to-one, setting $C_1[i + 1] \sim C_2[j + 1]$ (Fig. 13). By induction we can now compute correspondences for every point on the curves. This




Fig. 13. (a) Portions of two curves with possible pointwise correspondences shown (and labeled with graph notation, e.g., $v(i, j) \Leftrightarrow C_1[i] \sim C_2[j]$). (b) Typical portion of graph in Fig. 12 showing “current” correspondence node $v(i, j)$ and the three possible “next” nodes.

sketch of the algorithm lends itself to a graph formulation, where each node of a directed graph is a correspondence between a point on each curve, and each edge represents one of the possible actions, linking the current node to the possible “next” nodes. Assuming C_1 and C_2 are sampled with N and M points, respectively, the graph is a regular grid with NM nodes and $3NM$ edges. Because the curves are defined on a periodic domain, the graph has edges that “wrap around” from the top of the grid to the bottom, and from the right side of the grid to the left, as shown in Fig. 12.⁶ The edges are weighted by the distance between the invariants associated with the “next” node, cf., (15). We formalize and extend this concept in the remainder of this section.

We sample the curves C_1, C_2 on the discrete domains Ω_1, Ω_2 with equal spacing as follows:

$$\begin{aligned} \Omega_1 &= [0, \Delta s_1, 2\Delta s_1, \dots, M\Delta s_1 = 1], \\ \Delta s_1 &= \frac{1}{M}, \quad M \in \mathbb{N}^+, \end{aligned} \tag{24}$$

$$\begin{aligned} \Omega_2 &= [0, \Delta s_2, 2\Delta s_2, \dots, N\Delta s_2 = 1], \\ \Delta s_2 &= \frac{1}{N}, \quad N \in \mathbb{N}^+. \end{aligned} \tag{25}$$

The weighted, directed graph $G = (V, E)$ is formed based on a grid structure of the discrete domain $\Omega = \Omega_1 \times \Omega_2$, as shown in Fig. 12. Each node $v(i, j) \in V$ in the graph represents a pointwise correspondence $C_1(i\Delta s_1) \sim C_2(j\Delta s_2)$, where $i \in [0, M] \subset \mathbb{N}^+$ and $j \in [0, N] \subset \mathbb{N}^+$. The adjacency relation of nodes is defined by an edge $e(v(i, j), v(k, l))$ that represents a directed relation $v(i, j) \rightarrow v(k, l)$ indicating the following correspondence $v(k, l)$ given the current correspondence $v(i, j)$.

The minimization of the energy functional \tilde{E} is equivalent to finding a shortest path $p = \langle v_0, v_1, v_2, \dots, v_L \rangle$ that gives a minimum weight from $v_0 = v(0, 0)$ to $v_L = v(M, N)$,

$$w(p) = \sum_{t=0}^{L-1} w(v_t, v_{t+1}) \tag{26}$$

6. The graph domain can also be considered to be embedded in a two-dimensional torus.

based on a weighting function $w(v_t, v_{t+1})$ adapted from (15), defined by

$$\begin{aligned} w(v(i, j), v(k, l)) &= \|I_1(k\Delta s_1) - I_2(l\Delta s_2)\|^2 \frac{(h'_1 + h'_2)}{2} \\ &+ \alpha \left\| \frac{h'_2 - h'_1}{h'_1 + h'_2} \right\|^2 \frac{(h'_1 + h'_2)}{2}. \end{aligned} \tag{27}$$

The regular nature of the graph allows us to simplify the computations of $h'_{1,2}$:

$$\begin{aligned} h'_1 &= \frac{\Delta s_1}{\Delta s_1} = 1, & h'_2 &= \frac{0}{\Delta s_2} = 0, & \text{if } k &= \text{mod}(i, M) + 1, & l &= j, \\ h'_1 &= \frac{0}{\Delta s_1} = 0, & h'_2 &= \frac{\Delta s_2}{\Delta s_2} = 1, & \text{if } k &= i, & l &= \text{mod}(j, N) + 1, \\ h'_1 &= \frac{\Delta s_1}{\sqrt{\Delta s_1^2 + \Delta s_2^2}}, & h'_2 &= \frac{\Delta s_2}{\sqrt{\Delta s_1^2 + \Delta s_2^2}}, & \text{if } k &= \text{mod}(i, M) + 1, & l &= \text{mod}(j, N) + 1. \end{aligned} \tag{28}$$

The direction of edges in the graph is constrained so that the warping function $h(\xi)$ is monotonic. The monotonicity of the warping function prevents cross correspondence that causes a topological change in matching. (This is equivalent to

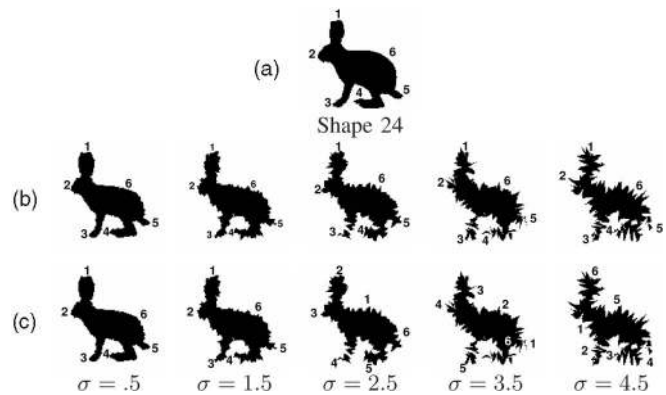


Fig. 14. Shape correspondence with increasing noise perturbation. (a) Correspondence of Shape 24 to noisy instances of Shape 20 computed, (b) via integral invariants, and (c) via differential invariants. Some salient points on the contour are labeled on Shape 24, and the corresponding points are labeled on the remaining contours. Since the integral invariant is more robust to noise than the differential one, it consistently identifies the corresponding parts, even for contours which are strongly perturbed by noise. For the differential invariant, on the other hand, the algorithm fails to capture the correct correspondence when $\sigma \geq 2.5$.




Fig. 15. Shape distance as a function of noise for the shapes in Fig. 14. While the shape distance measure based on differential invariants strongly varies with noise, the distance based on integral invariants is much more insensitive to noise. (See text for details.)

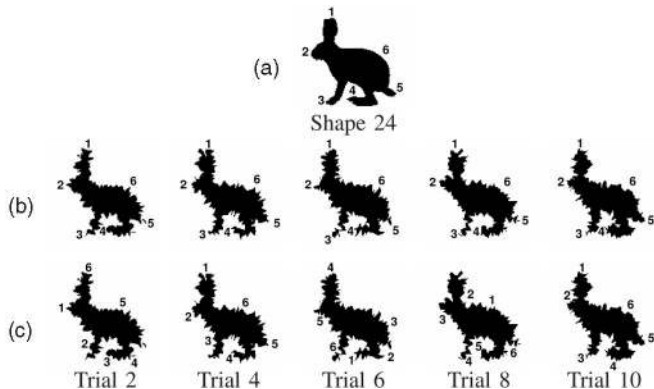


Fig. 16. Shape correspondence for several perturbations (with noise of scale $\sigma = 2.5$). (a) Correspondence of Shape 24 to noisy instances of Shape 20 computed, (b) via integral invariants and (c) via differential invariants. Some salient points on the contour are labeled on Shape 24, and the corresponding points are labeled on the remaining contours. In contrast to the distance based on differential invariants, the integral invariant distance consistently provides the correct correspondence.

enforcing the constraint in Section 4 that $-1 \leq d(s) \leq 1$. Further, each point in each curve must have at least one corresponding point in the other curve. These constraints are implied by the existence of an edge between nodes $v(i, j)$ and $v(k, l)$ only if $i \in \{\text{mod}(k, M) + 1\}$ and $j \in \{l, \text{mod}(l, N) + 1\}$.

Dijkstra's algorithm is used for finding a shortest path from a single source node to a single destination node in a graph. Let p be a sequence for the shortest path in the graph $G = (V, E)$ such as

$$p = \langle v_0, v_1, v_2, \dots, v_P \rangle \tag{29}$$

$$= \langle v(i_0, j_0), v(i_1, j_1), v(i_2, j_2), \dots, v(i_P, j_P) \rangle,$$

where P is the number of nodes in the path p , and $v_0 = v_P$. Then, the optimal warping function $h(\xi) = (h_1(\xi), h_2(\xi))$ is given by

$$h_1(\xi) = (i_0 \Delta s_1, i_1 \Delta s_1, i_2 \Delta s_1, \dots, i_{P-1} \Delta s_1) \tag{30}$$

$$h_2(\xi) = (j_0 \Delta s_2, j_1 \Delta s_2, j_2 \Delta s_2, \dots, j_{P-1} \Delta s_2).$$

An example of the result of this algorithm is shown in Fig. 11. Finally, the shape distance is computed via (26).




Fig. 17. Histograms of shape distance between Shape 24 and 1,000 perturbations of Shape 20 with noise at variance $\sigma = 2.5$. The dark bars on the left represent shape distances computed with integral invariants; the lighter bars on the right represent shape distances computed via differential invariants. At a fixed noise variance, the computed shape distance based on integral invariants remains essentially constant over all trials, in contrast to the distance based on differential invariants. (To scale the histograms for presentation, two experiments were removed from the differential invariant trials, with shape distances 2.8 and 9.7.)

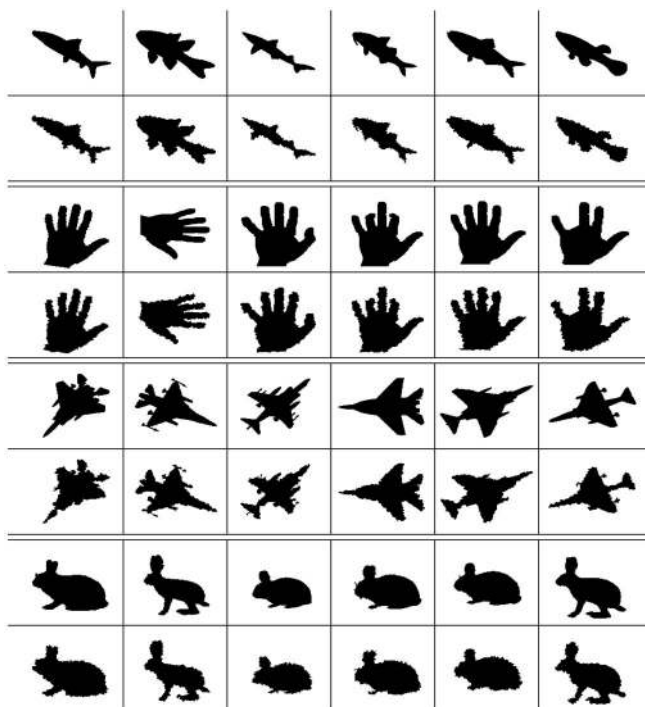


Fig. 18. Shapes used in matching experiments.

No fast algorithm exists to determine the best choice of the initial correspondence. Previous implementations (cited above) choose a fixed a point on the first curve and pair it with all possible choices of points on the second curve, calculating the path for each pair to determine the shortest. (This process can be thought of as exhaustively searching among all possible values of $d(0)$ or $h_1(0)$.) Searching near an initial point estimated as in [51] will obviate the need for an exhaustive search. Alternately, the exhaustive search can be




Fig. 20. Shape distance between noisy shapes (across top) and original shapes (along left side) via integral invariant. Lighter shade indicates higher distance. See text for more details.

result in a shape description that is less sensitive to occlusions or localized deformations, when compared to similarly implemented differential invariant methods. We compute the differential invariant using the method outlined in [13], [8]. We will begin with experiments demonstrating the computation of shape distance and correspondences between two shapes before demonstrating the retrieval of matches for noisy shapes from a database [76].

In this section, we use a subset of shapes from the Kimia silhouette database.⁷ The contours are extracted from the silhouettes using the pixel coordinates as a reference frame. The result is that the x and y -values of the contours are in $[0, 99]$. Empirically, we sample the contours at 100 evenly spaced points, and set $r = 15$ and $\alpha = .01$.

Fig. 14 shows the shape matching induced via the local-area integral invariant and via curvature between two different bunnies⁸ despite increasing noise. Noise is added by perturbing all points on the contour in the normal direction by a distance drawn from a zero-mean Gaussian random variable with specified σ . We indicate the correspondence by showing the mapping of the numbered landmarks onto the noisy shape, although we emphasize that invariant values from *everywhere* on the curve, and not just at feature points, are used to compute shape matching and distance. Fig. 15 is a plot of the shape distance for the matching, shown in Fig. 14.

7. Available at <http://www.lcms.brown.edu/vision/software/index.html>.

8. Although the two bunnies look similar, closer examination shows that the noisy bunny has a thicker body and a longer snout, in addition to differences in position.

Note that the distance computed via the integral invariant increases as σ increases, but, in general, added noise affects the shape distance only slightly. Contrast this with the distance computed via curvature, which increases drastically as a function of σ until the curves are so noisy that a meaningful correspondence cannot consistently be computed using differential invariants (e.g., the $\sigma = 2.5$ column of Fig. 14). The drastic *decrease* in shape distance for values of σ beyond this “breakdown value” further demonstrates the dangers of relying on differential invariants; even though the distance value indicates that this correspondence is optimal, the correspondence is subjectively incorrect.

Fig. 16 again shows the noise robustness of the integral invariant, compared to differential invariants, when used for shape matching. The variance of the noise is held constant at $\sigma = 2.5$, however, the experiment is repeated 1,000 times. Shape matching via the integral invariant provides a consistent correspondence (as shown with the labeled features in the second row) and a consistent shape distance shown in the histogram in Fig. 17. Computation of shape matching and shape distance via curvature results in a correspondence that varies with the noise, as shown in the third row, and a more erratic shape distance, shown in Fig. 17.

In Fig. 19, the results of matching and retrieving noisy shapes (shown on the left side) from a database (shown across the top) are shown. (Large displays of the original and noisy shapes are shown in Fig. 18.) We especially highlight several pairs where representation by differential invariants leads to mismatches, such as the third, fourth, and fifth fish (in the first group). Due to the differential invariant’s sensitivity to




Fig. 21. Shape distance between noisy shapes (across top) and original shapes (along left side) via differential invariant. Lighter shade indicates higher distance.

noise, these fish have a smaller shape distance to the rabbits, where the shape distance based on integral invariants orders the shapes correctly (fish are closer to themselves than to rabbits). Examination of the data shows several such cases where integral invariants are more robust than curvature on noisy shapes.

Fig. 20 and Fig. 21 show this same data in a more aggregate way, using shades of gray to indicate distance. The distance matrix computed using integral invariants (Fig. 20) has low distances on the diagonal and the block-diagonal structure, as expected in a database with grouped classes. Contrast this with the curvature-based distance matrix (Fig. 21) which lacks clearly lower distance on the diagonal and has a vertically *banded* structure opposed to the desired block-diagonal structure, indicating that the added noise influenced the shape distance more than the shape itself.

7 DISCUSSION AND CONCLUSIONS

In this paper, we address one of the key disadvantages of differential invariants for shape matching—namely, their inherent sensitivity to noise. We introduce a new distance for 2D shapes which is based on the notion of *integral* group-invariant descriptions of shape. Both theoretically and experimentally we relate these integral invariant shape distances to previously proposed shape distances which are based on differential invariants.

While integral invariants are employed for robustness to high-frequency noise and small deformations, shape warping by the computation of an optimal reparameterization allows us to account for large localized changes such as

occlusions and configuration changes. We embed both of these concepts in a formulation of a shape distance, and outline how distance and optimal correspondence are jointly computed via efficient dynamic programming algorithms.

On an experimental level, we demonstrate robustness of the integral invariant distances for shape matching and identifying corresponding shape parts under perturbation in increasing amounts of noise.

ACKNOWLEDGMENTS

This work was performed under the auspices of the US Department of Energy by the University of California Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48. This research was supported by grants NIH U54 RR021813, AFOSR F49620-03-1-0095/E-16-V91-G2, and ONR N00014-03-1-0850/N00014-02-1-0720.

REFERENCES

- [1] R. Alferéz and Y.F. Wang, "Geometric and Illumination Invariants for Object Recognition," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 21, no. 6, pp. 505-536, June 1999.
- [2] K. Arbter, W.E. Snyder, H. Burkhardt, and G. Hirzinger, "Applications of Affine-Invariant Fourier Descriptors to Recognition of 3-D Objects," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 12, no. 7, pp. 640-646, July 1990.
- [3] N. Ayache and O. 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.
- [4] M. Bakircioglu, U. Grenander, N. Khaneja, and M.I. Miller, "Curve Matching on Brain Surfaces Using Frenet Distances," *Human Brain Mapping*, vol. 6, pp. 329-333, 1998.

- [5] R. Basri, L. Costa, D. Geiger, and D. Jacobs, "Determining the Similarity of Deformable Shapes," *Vision Research*, vol. 38, pp. 2365-2385, 1998.
- [6] S. Belongie, J. Malik, and J. Puzicha, "Shape Matching and Object Recognition Using Shape Contexts," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 24, no. 4, pp. 509-522, Apr. 2002.
- [7] A. Bengtsson and J.-O. Eklundh, "Shape Representation by Multiscale Contour Approximation," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 13, no. 1, pp. 85-93, Jan. 1991.
- [8] M. Boutin, "Numerically Invariant Signature Curves," *Int'l J. Computer Vision*, vol. 40, no. 3, pp. 235-248, 2000.
- [9] R.D. Brandt and F. Lin, "Representations that Uniquely Characterize Images Modulo Translation, Rotation and Scaling," *Pattern Recognition Letters*, vol. 17, pp. 1001-1015, 1996.
- [10] A. Bruckstein, N. Katzir, M. Lindenbaum, and M. Porat, "Similarity Invariant Signatures for Partially Occluded Planar Shapes," *Int'l J. Computer Vision*, vol. 7, no. 3, pp. 271-285, 1992.
- [11] A.M. Bruckstein, R.J. Holt, A.N. Netravali, and T.J. Richardson, "Invariant Signatures for Planar Shape Recognition under Partial Occlusion," *J. Computer Vision, Graphics, and Image Processing*, vol. 58, no. 1, pp. 49-65, 1993.
- [12] A.M. Bruckstein, E. Rivlin, and I. Weiss, "Scale-Space Semi-Local Invariants," *Image and Vision Computing*, vol. 15, no. 5, pp. 335-344, 1997.
- [13] E. Calabi, P. Olver, C. Shakiban, A. Tannenbaum, and S. Haker, "Differential and Numerically Invariant Signature Curves Applied to Object Recognition," *Int'l J. Computer Vision*, vol. 26, pp. 107-135, 1998.
- [14] D. Chetverikov and Y. Khenokh, "Matching for Shape Defect Detection," *Lecture Notes in Computer Science*, vol. 1689, no. 2, pp. 367-374, 1999.
- [15] I. Cohen, N. Ayache, and P. Sulger, "Tracking Points on Deformable Objects Using Curvature Information," *Proc. European Conf. Computer Vision*, pp. 458-466, 1992.
- [16] T. Cohignac, C. Lopez, and J.M. Morel, "Integral and Local Affine Invariant Parameter and Application to Shape Recognition," *Proc. Int'l Conf. Pattern Recognition*, vol. 1, pp. 164-168, 1994.
- [17] J.B. Cole, H. Murase, and S. Naito, "A Lie Group Theoretical Approach to the Invariance Problem in Feature Extraction and Object Recognition," *Pattern Recognition Letters*, vol. 12, pp. 519-523, 1991.
- [18] D. Cremers, T. Kohlberger, and C. Schnörr, "Shape Statistics in Kernel Space for Variational Image Segmentation," *Pattern Recognition*, vol. 36, no. 9, pp. 1929-1943, 2003.
- [19] D. Cremers, S.J. Osher, and S. Soatto, "Kernel Density Estimation and Intrinsic Alignment for Knowledge-Driven Segmentation: Teaching Level Sets to Walk," *Pattern Recognition*, Sept. 2004.
- [20] D. Cremers and S. Soatto, "A Pseudo-Distance for Shape Priors in Level Set Segmentation," *Proc. IEEE Second Int'l Workshop Variational, Geometric, and Level Set Methods*, pp. 169-176, 2003.
- [21] R. Davies, C. Twining, T. Cootes, J. Waterton, and C Taylor, "A Minimum Description Length Approach to Statistical Shape Modeling," *IEEE Trans. Medical Imaging*, vol. 21, no. 5, pp. 525-537, 2002.
- [22] A. DelBimbo and P. Pala, "Visual Image Retrieval by Elastic Matching of User Sketches," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 19, no. 2, pp. 121-132, Feb. 1997.
- [23] A. Dervieux and F. Thomasset, "A Finite Element Method for the Simulation of Raleigh-Taylor Instability," *Springer Lecture Notes in Math.*, vol. 771, pp. 145-158, 1979.
- [24] L.E. Dickson, *Algebraic Invariants*. John Wiley and Sons, 1914.
- [25] J. Dieudonne and J. Carrell, *Invariant Theory: Old and New*. Academic Press, 1970.
- [26] I.L. Dryden and K.V. Mardia, *Statistical Shape Analysis*. Wiley, 1998.
- [27] J. Flusser and T. Suk, "Pattern Recognition by Affine Moment Invariants," *Pattern Recognition*, vol. 26, no. 1, pp. 167-174, 1993.
- [28] D.A. Forsyth, J.L. Mundy, A. Zisserman, and C.M. Brown, "Projectively Invariant Representations Using Implicit Algebraic Curves," *Image and Vision Computing*, vol. 9, no. 2, pp. 130-136, 1991.
- [29] D.A. Forsyth, J.L. Mundy, A.P. Zisserman, C. Coelho, A. Heller, and C.A. Othwell, "Invariant Descriptors for 3-D Object Recognition and Pose," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 13, no. 10, pp. 971-991, Oct. 1991.
- [30] Y. Gdalyahu and D. Weinshall, "Flexible Syntactical 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, Dec. 1999.
- [31] L. Van Gool, T. Moons, E. Pauwels, and A. Oosterlinck, "Semi-Differential Invariants," *Geometric Invariance in Computer Vision*, pp. 193-214, 1992.
- [32] L. Van Gool, T. Moons, and D. Ungureanu, "Affine/Photometric Invariants for Planar Intensity Patterns," *Proc. European Conf. Computer Vision*, vol. 1, pp. 642-651, 1996.
- [33] J.H. Grace and A. Young, *The Algebra of Invariants*. Cambridge Univ. Press, 1903.
- [34] C.E. Hann and M.S. Hickman, "Projective Curvature and Integral Invariants," *Int'l J. Computer Vision*, vol. 40, no. 3, pp. 235-248, 2000.
- [35] M.K. Hu, "Visual Pattern Recognition by Moment Invariants," *IRE Trans. IT*, vol. 8, pp. 179-187, 1961.
- [36] K. Kanatani, *Group Theoretical Methods in Image Understanding*. Springer, 1990.
- [37] E. Klassen, A. Srivastava, M. Mio, and S.H. Joshi, "Analysis of Planar Shapes Using Geodesic Paths on Shape Spaces," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 26, no. 3, pp. 372-383, Mar. 2004.
- [38] P.N. Klein, S. Tirthapura, D. Sharvit, and B. Kimia, "A Tree-Edit-Distance Algorithm for Comparing Simple, Closed Shapes," *Proc. Symp. Discrete Algorithms*, pp. 696-704, 2000.
- [39] E.P. Lane, *Projective Differential Geometry of Curves and Surfaces*. Univ. of Chicago Press, 1932.
- [40] J. Lasenby, E. Bayro-Corrochano, A.N. Lasenby, and G. Sommer, "A New Framework for the Formation of Invariants and Multiple-View Constraints in Computer Vision," *Proc. Int'l Conf. Image Processing*, pp. 313-316, 1996.
- [41] L.J. Latecki and R. Lakämper, "Shape Similarity Measure Based on Correspondence of Visual Parts," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 22, no. 10, pp. 1185-1190, Oct. 2000.
- [42] H. Le and D.G. Kendall, "The Riemannian Structure of Euclidean Shape Spaces: A Novel Environment for Statistics," *Annals of Statistics*, vol. 21, no. 3, pp. 1225-1271, 1993.
- [43] G. Lei, "Recognition of Planar Objects in 3-D Space from Single Perspective Views Using Cross Ratio," *Robotics and Automation*, vol. 6, no. 4, pp. 432-437, 1990.
- [44] R. Lenz, "Group Theoretical Methods in Image Processing," *Lecture Notes in Computer Science*, vol. 413, 1990.
- [45] M.E. Leventon, W.E.L. Grimson, and O. Faugeras, "Statistical Shape Influence in Geodesic Active Contours," *Proc. Conf. Computer Vision and Pattern Recognition*, vol. 1, pp. 316-323, June 2000.
- [46] S.Z. Li, "Shape Matching Based on Invariants," *Progress in Neural Networks: Shape Recognition*, vol. 6, pp. 203-228, 1999.
- [47] S. Liao and M. Pawlak, "On Image Analysis by Moments," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 18, no. 3, pp. 254-266, Mar. 1996.
- [48] H. Liu and M. Srinath, "Partial Shape Classification Using Contour Matching in Distance Transforms," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 12, no. 2, pp. 1072-1079, Feb. 1990.
- [49] T. Liu and D. Geiger, "Approximate Tree Matching and Shape Similarity," *Proc. Int'l Conf. Computer Vision*, pp. 456-462, 1999.
- [50] S. Manay, B. Hong, A. Yezzi, and S. Soatto, "Integral Invariant Signatures," *Proc. European Conf. Computer Vision*, May 2004.
- [51] J.S. Marques and A.J. Abrantes, "Shape Alignment—Optimal Initial Point and Pose Estimation," *Pattern Recognition Letters*, vol. 18, no. 1, pp. 49-53, 1997.
- [52] T. Miyatake, T. Matsuyama, and M. Nagao, "Affine Transform Invariant Curve Recognition Using Fourier Descriptors," *Information Processing Soc. Japan*, vol. 24, no. 1, pp. 64-71, 1983.
- [53] F. Mokhtarian and A.K. Mackworth, "Scale-Based Description and Recognition of Planar Curves and Two-Dimensional Shapes," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 8, no. 1, pp. 34-43, Jan. 1986.
- [54] F. Mokhtarian and A.K. Mackworth, "A Theory of MultiScale, Curvature-Based Shape Representation for Planar Curves," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 14, no. 8, pp. 789-805, Aug. 1992.
- [55] D. Mumford, "Mathematical Theories of Shape: Do They Model Perception," *Geometric Methods in Computer Vision*, vol. 1570, pp. 2-10, 1991.
- [56] D. Mumford, J. Fogarty, and F.C. Kirwan, *Geometric Invariant Theory*, third ed. Springer-Verlag, 1994.
- [57] D. Mumford, A. Latto, and J. Shah, "The Representation of Shape," *Proc. IEEE Workshop Computer Vision*, pp. 183-191, 1984.
- [58] *Geometric Invariance in Computer Vision*. J.L. Mundy and A. Zisserman, eds. MIT Press, 1992.

- [59] L. Nielsen and G. Saprre, "Projective Area-Invariants as an Extension of the Cross-Ratio," *J. Computer Vision, Graphics, and Image Processing*, vol. 54, no. 1, pp. 145-159, 1991.
- [60] P.J. Olver, *Equivalence, Invariants and Symmetry*. Cambridge, 1995.
- [61] S.J. Osher and J.A. Sethian, "Fronts Propagation with Curvature Dependent Speed: Algorithms Based on Hamilton-Jacobi Formulations," *J. Computer Physics*, vol. 79, pp. 12-49, 1988.
- [62] T. Pajdla and L. Van Gool, "Matching of 3-D Curves Using Semi-Differential Invariants," *Proc. Int'l Conf. Computer Vision*, pp. 390-395, 1995.
- [63] M. Pelillo, K. Siddiqi, and S.W. Zucker, "Matching Hierarchical Structures Using Association Graphs," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 21, no. 11, pp. 1105-1120, Nov. 1999.
- [64] A. Pikaz and I. Dinstein, "Matching of Partially Occluded Planar Curves," *Pattern Recognition*, vol. 28, no. 2, pp. 199-209, 1995.
- [65] A. Pitiot, H. Delingette, A. Toga, and P. Thompson, "Learning Object Correspondences with the Observed Transport Shape Measure," *Proc. Conf. Information Processing in Medical Imaging*, 2003.
- [66] T.H. Reiss, "Recognizing Planar Objects Using Invariant Image Features," *Lecture Notes in Computer Science*, vol. 676, 1993.
- [67] C. Rothwell, A. Zisserman, D. Forsyth, and J. Mundy, "Canonical Frames for Planar Object Recognition," *Proc. European Conf. Computer Vision*, pp. 757-772, 1992.
- [68] C. Rothwell, A. Zisserman, D. Forsyth, and J. Mundy, "Planar Object Recognition Using Projective Shape Representation," *Int'l J. Computer Vision*, vol. 16, pp. 57-99, 1995.
- [69] M. Rousson and N. Paragios, "Shape Priors for Level Set Representations," *Proc. European Conf. Computer Vision*, May 2002.
- [70] 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.
- [71] G. Sapiro and A. Tannenbaum, "Affine Invariant Scale Space," *Int'l J. Computer Vision*, vol. 11, no. 1, pp. 25-44, 1993.
- [72] G. Sapiro and A. Tannenbaum, "Area and Length Preserving Geometric Invariant Scale-Spaces," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 17, no. 1, pp. 67-72, Jan. 1995.
- [73] J. Sato and R. Cipolla, "Affine Integral Invariants for Extracting Symmetry Axes," *Image and Vision Computing*, vol. 15, no. 8, pp. 627-635, 1997.
- [74] H. Schulz-Mirbach, "Invariant Features for Gray Scale Images," *Proc. 17th Ann. Symp. German Assoc. for Pattern Recognition*, pp. 1-14, 1995.
- [75] J. Schwartz and M. Sharir, "Identification of Partially Obscured Objects in Two and Three Dimensions by Matching Noisy Characteristic Curves," *Int'l J. Robotic Research*, vol. 6, no. 2, pp. 29-44, 1987.
- [76] T. Sebastian, P. Klein, and B. Kimia, "On Aligning Curves," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 25, no. 1, pp. 116-125, Jan. 2003.
- [77] E. Sharon and D. Mumford, "2D-Shape Analysis Using Conformal Mapping," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, June 2004.
- [78] D. Sharvit, J. Chan, H. Tek, and B. Kimia, "Symmetry-Based Indexing of Image Databases," *Proc. IEEE Workshop Content-Based Access of Image and Video Libraries*, pp. 56-62, 1998.
- [79] A. Shashua and N. Navab, "Relative Affine Structure: Canonical Model for 3D From 2D Geometry and Applications," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 18, no. 9, pp. 873-883, Sept. 1996.
- [80] K. Siddiqi, A. Shokoufandeh, S.J. Dickinson, and S.W. Zucker, "Shock Graphs and Shape Matching," *Proc. Int'l Conf. Computer Vision*, pp. 222-229, 1998.
- [81] C.E. Springer, *Geometry and Analysis of Projective Spaces*. San Francisco: Freeman, 1964.
- [82] H. Tagare, D. O'Shea, and A. Rangarajan, "A Geometric Correspondence for Shape-Based NonRigid Correspondence," *Proc. Int'l Conf. Computer Vision*, pp. 434-439, 1995.
- [83] Q.M. Tieng and W.W. Boles, "Recognition of 2D Object Contours Using the Wavelet Transform Zero-Crossing Representation," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 19, no. 8, pp. 910-916, Aug. 1997.
- [84] C. Tomasi and R. Manduchi, "Stereo without Search," *Proc. European Conf. Computer Vision*, pp. 452-465, 1996.
- [85] A. Trounev and L. Younes, "Diffeomorphic Matching Problems in One Dimension: Designing and Minimizing Matching Functions," *Proc. European Conf. Computer Vision*, pp. 573-587, 2000.
- [86] A. Tsai, A. Yezzi, W. Wells, C. Tempny, D. Tucker, A. Fan, E. Grimson, and A. Willsky, "Model-Based Curve Evolution Technique for Image Segmentation," *Proc. Conf. Computer Vision and Pattern Recognition*, pp. 463-468, 2001.
- [87] 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, Feb. 1993.
- [88] J. Verestoy and D. Chetverikov, "Shape Detection in Ferrite Cores," *Machine Graphics and Vision*, vol. 6, no. 2, pp. 225-236, 1997.
- [89] I. Weiss, "Noise Resistant Invariants of Curves," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 15, no. 9, pp. 943-948, Sept. 1993.
- [90] A.P. Witkin, "Scale-Space Filtering," *Proc. Int'l Joint Conf. Artificial Intelligence*, pp. 1019-1021, 1983.
- [91] H. Wolfson, "On Curve Matching," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 12, no. 5, pp. 483-489, May 1990.
- [92] L. Younes, "Optimal Matching between Shapes via Elastic Deformations," *Image and Vision Computing*, vol. 17, pp. 381-389, 1999.
- [93] C.T. Zahn and R.Z. Roskies, "Fourier Descriptors for Plane Closed Curves," *IEEE Trans. Computers*, vol. 21, pp. 269-281, 1972.
- [94] S. Zhu and A. Yuille, "Forms: A Flexible Object Recognition and Modeling System," *Int'l J. Computer Vision*, vol. 20, no. 3, pp. 187-212, 1996.
- [95] A. Zisserman, D.A. Forsyth, J.L. Mundy, C.A. Rothwell, and J.S. Liu, "3D Object Recognition Using Invariance," *Artificial Intelligence*, vol. 78, pp. 239-288, 1995.



lar focus on models and invariants of shape.



Research in Princeton, New Jersey. Since October 2005, he has headed the Research Group for Computer Vision and Pattern Recognition at the University of Bonn, Germany. His research is focused on statistical and variational methods for computer vision. He received several awards, among them the Best Paper of the Year 2003 by the Pattern Recognition Society, the Olympus Award 2004, and the UCLA Chancellor's Award for Postdoctoral Research 2005. He is a member of the IEEE.



registration, segmentation, and medical imaging.

Siddharth Manay received the PhD degree in 2003 from the School of Electrical Engineering, Georgia Institute of Technology, Atlanta. In 2003, he spent a year as a postdoctoral fellow with the Vision Lab at the Computer Science Department, University of California at Los Angeles. He is currently with Lawrence Livermore National Laboratory in Livermore, California. His research interests are in the field image processing and computer vision, with a particular

Daniel Cremers received the BS degree in mathematics (1994) and physics (1994), and the MS degree (Diplom) in theoretical physics (1997) from the University of Heidelberg. In 2002, he obtained the PhD degree in computer science from the University of Mannheim, Germany. Subsequently, he spent two years as a postdoctoral researcher at the University of California at Los Angeles and one year as a permanent researcher at Siemens Corporate

Byung-Woo Hong received the BSc degree in mathematics from Chung-Ang University, Korea in 1995, the MSc degree in computer science from the Weizmann Institute of Science, Israel in 2001, and the DPhil degree in information engineering from the University of Oxford, United Kingdom in 2005. He joined the computer vision lab at the University of California at Los Angeles as a postdoctoral researcher in 2005. His research interests include shape analysis,



Anthony J. Yezzi Jr. received the PhD degree in 1997 through the Department of Electrical Engineering at the University of Minnesota. After completing a postdoctoral research position in the Laboratory for Information and Decision Systems (LIDS) at the Massachusetts Institute of Technology, he joined the faculty of the School of Electrical and Computer Engineering at Georgia Institute of Technology in 1999 where he currently holds the position of associate

professor. Dr. Yezzi has also consulted for a number of medical imaging companies including GE, Picker, and VTI, and has been an IEEE member since 1999. His research lies primarily within the fields of image processing and computer vision. He has worked on a variety of problems including image denoising, edge-detection, segmentation and grouping, shape analysis, multiframe stereo reconstruction, tracking, and registration. Some central themes of his research include curve and surface evolution theory, differential geometry, and partial differential equations.



Stefano Soatto received the PhD degree in control and dynamical systems from the California Institute of Technology in 1996. He currently is a professor of computer science at the University of California at Los Angeles (UCLA). He joined UCLA after being an assistant and then associate professor of electrical and biomedical engineering at Washington University, and a research associate in applied sciences at Harvard University. He was also Ricercatore in the Department of Mathematics and Computer Science at the University of Udine, Italy. He received the DIng degree (highest honors) from the University of Padova, Italy in 1992. His general research interests are in computer vision and nonlinear estimation and control theory. In particular, he is interested in ways for computers to use sensory information (e.g., vision, sound, touch) to interact with humans and the environment. Dr. Soatto is the recipient of the David Marr Prize for work on Euclidean reconstruction and reprojection up to subgroups. He also received the Siemens Prize with the Outstanding Paper Award from the IEEE Computer Society for work on optimal structure from motion. He received the US National Science Foundation Career Award and the Okawa Foundation Grant. He is an associate editor of the *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, and a member of the editorial board of the *International Journal of Computer Vision (IJCV)* and of the journal *Foundations and Trends in Computer Graphics and Vision*. He is a member of the IEEE.

▷ **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.**