Integral Invariant Signatures

Siddharth Manay[1], Byung-Woo Hong[2], Anthony J. Yezzi[3], and Stefano Soatto[1,⋆]

[1] University of California at Los Angeles, Los Angeles CA 90024, USA,
[2] University of Oxford, Oxford OX1 3BW, UK
[3] Georgia Institute of Technology, Atlanta GA 30332, USA
{manay,hong,soatto}@cs.ucla.edu ayezzi@ece.gatech.edu

Abstract. For shapes represented as closed planar contours, we introduce a class of functionals that are invariant with respect to the Euclidean and similarity group, obtained by performing integral operations. While such integral invariants enjoy some of the desirable properties of their differential cousins, such as locality of computation (which allows matching under occlusions) and uniqueness of representation (in the limit), they are not as sensitive to noise in the data. We exploit the integral invariants to define a unique signature, from which the original shape can be reconstructed uniquely up to the symmetry group, and a notion of scale-space that allows analysis at multiple levels of resolution. The invariant signature can be used as a basis to define various notions of distance between shapes, and we illustrate the potential of the integral invariant representation for shape matching on real and synthetic data.

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, 3-D 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 millennium, 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

⋆ Supported by NSF IIS-0208197, AFOSR F49620-03-1-0095, ONR N00014-03-1-0850

T. Pajdla and J. Matas (Eds.): ECCV 2004, LNCS 3024, pp. 87–99, 2004.

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.

In this paper, we introduce a general class of invariants that are *integral* functionals of the data, as opposed to differential ones. We argue that such functionals are far less sensitive to noise, while retaining the nice features of differential invariants such as locality, which allow for matching under occlusions. They can be exploited to define invariant signature curves that can be used as a representation to define various notions of distances between shapes. We restrict our analysis to Euclidean and similarity invariants, although extensions to the affine group are straightforward. The integration kernel allows us to define intrinsic scale-spaces of invariant signatures, so that we can represent shapes at different levels of resolution and under various levels of measurement noise. We also show that our invariants can be computed very efficiently without performing explicit sums (in the discretized domain). Finally, we show that in the limit where the kernel measure goes to zero, one class of integral invariant is in one-to-one correspondence with the prince of differential invariants, curvature. This allows the establishment of a completeness property of the representation, in the limit, in that a given shape can be reconstructed uniquely, up to the invariance group, from its invariant signature. This relationship allows us to tap into the rich literature on differential invariants for theoretical results, while in our experiments we can avoid computing higher-order derivatives. We illustrate our results with several experiments, showed as space allows.

2 Relation to Existing Work, and Our Contribution

The role of invariants in computer vision has been advocated for various applications ranging from shape representation [34,4] to shape matching [3,29], quality control [48,11], and general object recognition [39,1]. Consequently a number of features that are invariant under specific transformations have been investigated [14,25,15,21,33,46]. In particular, one can construct primitive invariants of algebraic entities such as lines, conics and polynomial curves, based on a global descriptor of shape [36,18]. 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; while generic viewpoint invariants do not exist, invariant features can be computed from a collection of coplanar points or lines [40,41,20,6,17,52,1,45,26]. An invariant descriptor of a collection of points that relates to our approach is the shape context introduced by Belongie et al. [3], which consists in 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 [28,24,13,35]. 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 [10,37]. 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 [5].

In order to reduce noise-induced fluctuations of the signature, semi-differential invariants methods are introduced by using first derivatives and one reference point instead of curvature, thus avoiding the computation of high-order derivatives [38,19,27]. Another semi-invariant is given by transforming the given coordinate system to a canonical one [49].

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 [8,7]. 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 [50]. The combination of invariant theory with geometric multiscale analysis is investigated by applying an invariant diffusion equation for curve evolution [42,43,12]. A scale-space can be determined by varying the size of the differencing interval used to approximate derivatives using finite differences [9]. In [32], 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 [23]. Moment invariants under affine transformations were derived from the classical moment invariants in [16]. 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 [30]. The Fourier transform was also applied to obtain integral invariants [51,31,2]. 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 [47] and potentials were also proposed to preserve locality [22]. Alternatively, semi-local integral invariants are presented by integrating object curves with respect to arc length [44].

In this manuscript, we introduce two general classes of integral invariants; for one of them, we show its relationship to differential invariants (in the limit), which allows us to conclude that the invariant signature curve obtained from the integral invariant is in one-to-one correspondence with the original shape, up to the action of the nuisance group. We use the invariant signature to define various notions of distance between shapes, and we illustrate the potential of our representation on several experiments with real and simulated images.

3 Integral Invariants

Throughout this section we indicate with $\gamma : \mathbb{S}^1 \to \mathbb{R}^2$ a closed planar contour with 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 $\bar{\gamma}$ to indicate either the interior of the region bounded by γ (a two-dimensional object), or the curve γ itself (a one-dimensional object), and $d\mu(x)$ the corresponding measure, i.e. the area form dx or the arclength $ds(x)$ respectively. With this notation, we can define a fairly general notion of integral invariant.

Definition 1. *A function $I_\gamma(p) : \mathbb{R}^2 \to \mathbb{R}$ is an integral G-invariant if there exists a kernel $h : \mathbb{R}^2 \times \mathbb{R}^2 \to \mathbb{R}$ such that*

$$I_\gamma(p) = \int_{\bar{\gamma}} h(p, x) d\mu(x) \tag{1}$$

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

$$\int_{\bar{\gamma}} h(p, x) d\mu(x) = \int_{g\bar{\gamma}} h(gp, x) d\mu(x) \;\forall\; g \in G. \tag{2}$$

where $g\gamma \doteq \{gx \mid g \in G, x \in \gamma\}$, and similarly for $g\bar{\gamma}$.

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

Example 1 (Integral distance invariant). *Consider $G = SE(2)$ and the following function, computed at every point $p \in \gamma$:*

$$I_\gamma(p) \doteq \int_\gamma d(p, x) ds(x) \tag{3}$$

where $d(x, y) \doteq |y - x|$ is the Euclidean distance in $\mathbb{R}^2$. This is illustrated in Fig. 1-a.

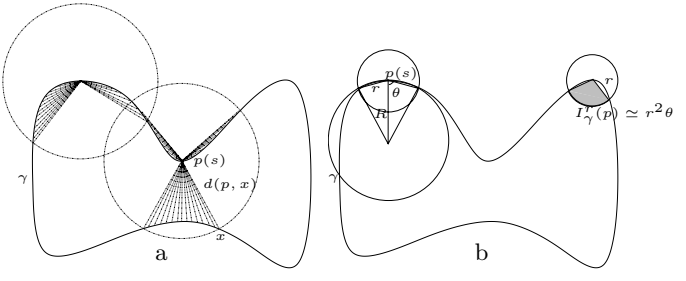

Fig. 1. (Left) Integral distance invariant defined in eq. (3), made local by means of a kernel as described in eq. (5). (Right) Integral area invariant defined by eq. (6).

It is immediate to show that this is an integral Euclidean invariant. The function I_γ associates to each point on the contour a number that is the average distance from that point to every other point on the contour. In particular, if the point $p \in \gamma$ 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 positive reals:

$$\{\gamma : \mathbb{S}^1 \to \mathbb{R}^2\} \mapsto \{I_\gamma(p(s)) : [0,\ L] \to \mathbb{R}_+.\} \tag{4}$$

This invariant is computed for a few representative shapes in Fig. 2 and Fig. 3.

A more "local" version of the invariant signature I_γ can be obtained by weighting the integral in eq. (3) with a kernel $q(p, x)$, so that $I_\gamma(p) \doteq \int_\gamma h(p, x)ds(x)$ where

$$h(p, x) \doteq q(p, x)d(p, x). \tag{5}$$

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 version of the "shape context," which was designed for a finite collection of points [3]. 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.

Example 2 (Integral area invariant). *Consider now the kernel $h(p, x) = \chi(B_r(p) \cap \bar{\gamma})(x)$, which represents the indicator function of the intersection of a small circle of radius r centered at the point p with the interior of the curve γ. For any given radius r, the corresponding integral invariant*

$$I^r_\gamma(p) \doteq \int_{B_r(p)\cap\bar{\gamma}} dx \tag{6}$$

can be thought of as a function from the interval $[0,\ L]$ to the positive reals, bounded above by the area of the region bounded by the curve γ. This is illustrated in Fig. 1-b and examples are shown in Fig. 2 and Fig. 3.

Naturally, if we plot the value of $I^r_\gamma(p(s))$ for all values of s and r ranging from zero to a maximum radius so that the local kernel encloses the entire curve $B_r(p) \supset \gamma$, we can generate a graph of a function that can be interpreted as a scale-space of integral invariants. Furthermore, $\chi(B_r(p))$ can be substituted by a more general kernel, for instance a Gaussian centered at p with $\sigma = r$.

Example 3 (Differential invariant). *Note that a regularized version of curvature, or in general a curvature scale space, can be interpreted as an integral invariant, since regularized curvature is an algebraic function of the first- and second-regularized derivatives [32]. Therefore, integral invariants are more general, but we will not exploit this added generality, since it contrary to the spirit of this manuscript, that is of avoiding the computation of derivatives of the image data, even if regularized.*

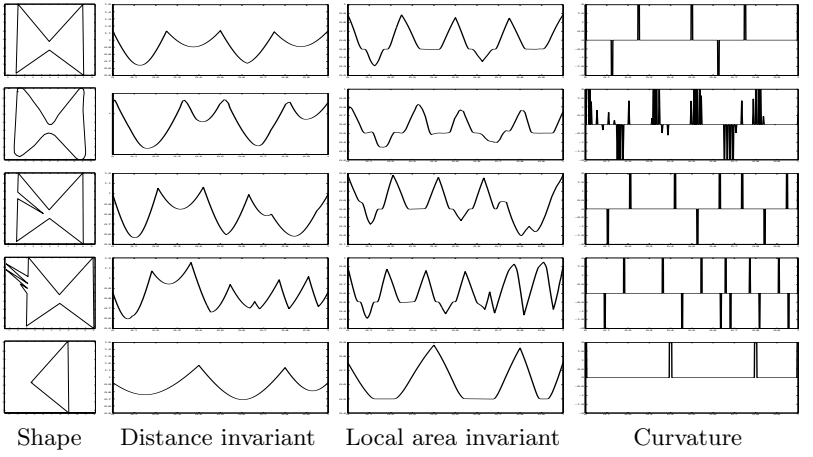

Fig. 2. For a set of representative shapes (left column), we compute the distance integral invariant of eq. (3) (middle left column), the local area invariant of eq. (6) with a kernel size $\sigma = 2$ (middle right column). Compare the results with curvature, shown in the rightmost column.

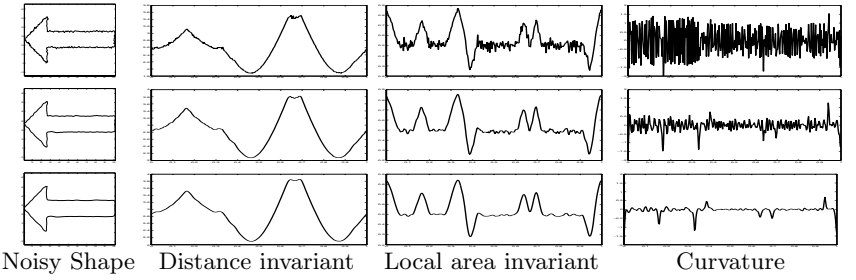

Fig. 3. For a noisy shape (left column), the distance invariant of eq. (3) with a kernel size of $\sigma = 30$ (middle left column), the local area invariant of eq. (6) with kernel size $r = 10$ (middle right column) and the differential invariant, curvature (right column). As one can see, noise is amplified in the computation of derivatives necessary to extract curvature.

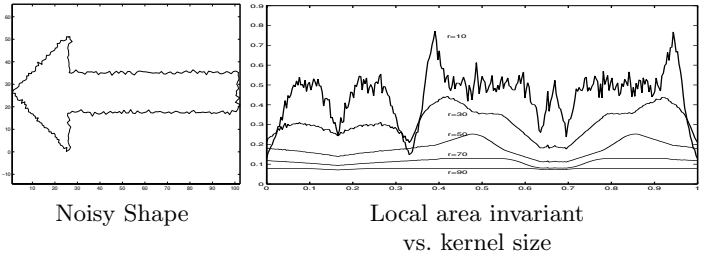

Fig. 4. For a noisy shape (left), the local area invariant of eq. (6) as a function of kernel size induces a scale-space of responses.

4 Relationship with Curvature and Local Differential Invariants

In this section we study the relationship between the local area invariant (6) and curvature. This is motivated by the fact that curvature is 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 [49], and therefore linking our integral invariant to curvature would allow us to tap onto the rich body of results on differential invariants without suffering from the shortcomings of computing high-order derivatives of the data.

We first assume that γ is smooth, so that a notion of curvature is well-defined, and the curve can be approximated locally by the osculating circle[1] $B_R(p)$ (Fig. 1-b). The invariant $I^r_\gamma(p)$ denotes the area of the intersection of a circle $B_r(p)$ with the interior of γ, and it can be approximated to first-order by the area of the shaded sector in Fig. 1-b, i.e. $I^r_\gamma(p) \simeq r^2\theta(p)$. Now, the angle θ can be computed as a function of r and R using the cosine law: $\cos\theta = r/2R$, and since curvature κ is the inverse of R we have

$$I^r_\gamma(p) \simeq r^2 \arccos\left(\frac{1}{2}r\kappa(p)\right). \tag{7}$$

Now, since arc-cosine 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 $r \to 0$; as r increases, $B_r(p)$ encloses the entire curve γ (which is closed), and consequently I^r_γ becomes a constant beyond a certain radius $r = r_{max}$. Therefore, for values of r that range from 0 to r_{max} we obtain an *intrinsic scale-space* of invariants, in contrast to the extrinsic scale-space of curvature. We compare these two descriptors in Fig. 3 and Fig. 4.

Note also that the integral invariant can be normalized via $I^r_\gamma/\pi r^2$ so as to provide a *scale-invariant* description of the curve, which is therefore invariant with respect to the similarity group. The corresponding integral invariant is then bounded between 0 and 1.

5 Invariant Signature Curves

The invariant $I^r_\gamma(p(s))$ can be represented by a function of s for any fixed value of r. This means, however, that in order to register two shapes, an "initial point" $s = 0$ must be chosen. There is nothing intrinsic to the geometry of the curve in the choice of this initial point, and indeed it would be desirable to devise a description that, in addition to being invariant to the group, is invariant with respect to the choice of initial point.

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

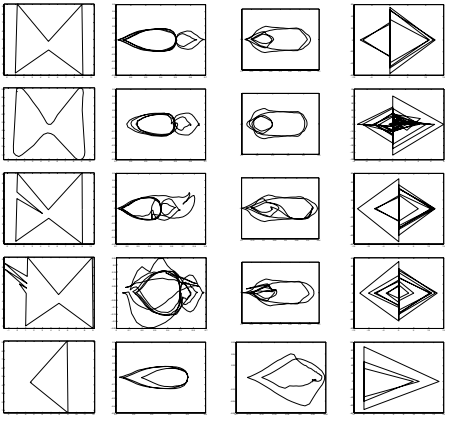

Fig. 5. Example of signature curves for a set of representative shapes (left column); local area invariant with small kernel (middle left column) and large kernel (middle right column), differential invariant (right column).

In order to do so, we follow the classic literature on differential invariants (see [10] and references therein) and plot a *signature*, that is the graph of $\frac{\partial I^r_\gamma(p(s))}{\partial s}$ versus I^r_γ. We indicate such a signature concisely by

$$(\dot{I}^r_\gamma, I^r_\gamma) \tag{8}$$

which of course can be plotted for all values of $r \in [0,\ r_{max}]$, yielding a scale-space of signatures. Naturally, we want to avoid direct computation of the derivative of the invariant, so the signature can be computed more simply as follows: Consider the binary image $\chi(\bar{\gamma})$ and convolve it with the kernel $h(p,x) \doteq B_r(p-x)$, where $p \in \mathbb{R}^2$, not just the curve γ. Evaluating the result of this convolution on $p \in \gamma$ yields I^r_γ, without the need to parameterize the curve. For $\dot{I}^r_\gamma$, compute the gradient of the filter response and inner-multiply the result with the tangent vector field of the image $\chi(\bar{\gamma})$, formed by filtering again by a kernel *different* than $B_r(p-x)$ and rotating its normalized gradient by 90^o. The result, when evaluated at $p \in \gamma$, yields $\dot{I}^r_\gamma$.

Notice that from the integral invariant signature we can reconstruct all differential invariants in the limit when $r \to 0$. In fact, from I^r_γ we can compute κ, and therefore from the signature we can compute $\dot{\kappa}$.

6 Distance between Shapes

In this section we outline methods for computing the distance between two shapes based on their invariants and invariant signatures curves.

	2	8	18	18	8	15	17	14	11	10	10	12	10	9	9	6	17	14	29	31	29	27	28
	40	78	78	77	70	359	466	428	322	337	*333*	378	369	373	386	383	361	393	541	526	557	484	484
	9	*1*	18	20	8	11	14	11	12	10	10	12	7	11	9	7	17	12	22	23	21	22	22
	67	*53*	77	80	72	377	482	441	330	349	386	386	389	383	400	399	373	406	548	535	567	489	496
	19	17	*2*	28	16	26	28	27	24	26	21	27	22	25	24	20	32	27	37	34	33	33	35
	70	84	*46*	84	72	378	487	443	333	356	401	393	388	389	401	406	382	413	557	540	580	498	500
	19	18	32	*2*	14	16	20	18	19	18	20	19	17	19	19	19	20	18	36	34	33	33	33
	75	89	90	*57*	80	381	511	468	356	371	417	409	402	408	420	419	400	429	567	546	589	503	513
	8	5	15	14	*1*	10	12	8	9	6	11	11	9	9	9	7	10	9	24	26	23	23	24
	56	69	67	71	*40*	332	439	*392*	279	311	346	*343*	343	341	358	*352*	328	360	*507*	*493*	*530*	*454*	458
	16	13	28	18	11	*0*	23	14	17	11	17	15	16	16	6	7	20	10	34	36	34	35	34
	197	212	226	224	197	*317*	668	596	493	508	556	535	546	541	540	363	521	572	763	732	789	689	681
	22	15	35	25	16	23	*1*	18	15	20	10	20	14	20	20	16	19	20	28	27	22	27	26
	244	247	246	264	240	567	*435*	663	522	548	613	588	579	588	578	579	595	622	815	775	826	751	732
	14	11	30	19	10	15	19	*0*	8	12	17	11	13	5	14	11	11	13	25	25	24	26	24
	208	223	231	240	210	539	678	397	486	521	557	534	547	358	564	558	547	551	769	749	801	708	721
	13	12	27	21	11	16	14	8	*0*	14	14	8	11	12	14	11	12	9	24	24	24	27	23
	197	216	215	217	186	519	662	570	*270*	474	545	526	543	523	542	544	521	559	748	736	785	685	701
	12	10	28	21	9	11	19	12	13	*1*	14	15	14	13	10	9	19	13	31	33	28	29	31
	221	231	236	237	212	507	686	601	469	*300*	557	566	560	547	567	560	498	577	771	750	794	717	716
	14	13	25	22	13	17	9	16	14	14	*1*	14	13	13	16	12	20	15	28	27	24	28	26
	205	229	251	248	218	546	688	608	513	535	350	557	544	566	593	545	549	595	789	774	803	727	713
	14	14	29	21	13	15	19	11	8	14	14	*1*	13	12	14	9	21	12	25	27	28	28	26
	224	240	240	251	226	555	683	643	520	529	575	357	563	571	582	568	558	599	794	762	800	729	725
	13	9	23	19	11	16	14	12	10	14	12	13	*1*	10	14	9	16	7	26	27	25	24	25
	222	239	232	244	225	559	680	631	501	510	581	557	*337*	534	580	576	541	374	773	740	798	720	698
	12	11	25	20	11	15	20	5	12	13	15	13	11	*0*	14	9	19	9	33	34	32	30	32
	194	208	214	223	180	520	646	394	456	496	526	536	532	*311*	553	548	519	542	738	710	765	663	675
	13	11	25	19	11	6	19	14	13	10	16	15	14	13	*1*	10	16	10	36	37	34	33	35
	223	245	248	249	218	550	675	629	506	501	584	561	571	553	*328*	582	548	602	789	767	810	724	736
	8	8	20	21	10	7	15	11	10	9	12	8	9	8	9	*1*	17	12	27	28	26	28	26
	213	224	231	232	207	345	676	613	471	514	541	544	556	545	564	355	532	573	783	743	807	712	708
	18	16	35	22	12	19	21	12	13	20	23	22	18	20	17	18	*0*	13	39	39	36	37	37
	191	198	211	218	184	529	659	607	467	492	536	526	540	528	538	519	*310*	538	751	723	775	688	689
	16	12	30	18	11	11	20	13	9	14	16	12	7	10	11	12	13	*0*	28	30	29	29	28
	180	200	210	216	178	515	642	598	470	502	542	528	344	480	539	518	506	*353*	740	707	765	668	675
	36	27	44	40	29	37	29	24	26	33	29	28	27	34	37	29	36	29	*1*	5	10	13	4
	314	332	338	334	325	696	831	811	666	689	726	727	734	715	693	722	699	751	522	910	972	864	868
	36	27	46	38	30	38	26	25	25	33	27	28	27	34	37	29	36	29	4	*1*	6	10	1
	248	262	259	262	258	607	753	703	569	607	662	634	640	644	642	619	628	646	824	514	843	767	*448*
	34	24	44	36	28	34	22	23	24	28	24	28	24	32	33	25	33	28	10	6	*1*	11	6
	290	310	311	319	303	682	798	796	631	675	737	677	707	708	714	720	687	719	902	873	559	802	825
	33	26	38	36	29	35	26	25	26	29	28	27	25	30	32	28	35	29	13	10	10	*1*	10
	296	298	308	304	303	671	789	751	622	635	697	698	677	660	700	689	650	705	861	847	910	477	836
	36	27	46	37	30	38	26	25	25	33	27	28	27	35	37	29	36	30	4	*1*	5	10	1
	246	259	259	262	259	596	773	713	586	593	650	653	643	644	661	619	612	637	830	514	852	765	448

Fig. 6. Noisy shape recognition from a database of 23 shapes. The upper number in each cell is the distance computed via the local-area integral invariant; the lower number is the distance computed via curvature invariant. The number in italics represents the best match for a noisy shape. See the text for more details

A straightforward distance between two shapes γ_1 and γ_2 is to compute a measure of the error between their invariants. One choice is the squared error.

$$D_E(\gamma_i, \gamma_j, r) = \int_0^1 (I^r_{\gamma_i}(p(s)) - I^r_{\gamma_j}(p(s)))^2 ds. \tag{9}$$

While this squared error can be computed for any invariant functional, we focus on invariants that preserve locality, such as the local area invariant, so that these distances will be valid for application to shape recognition despite occlusion.

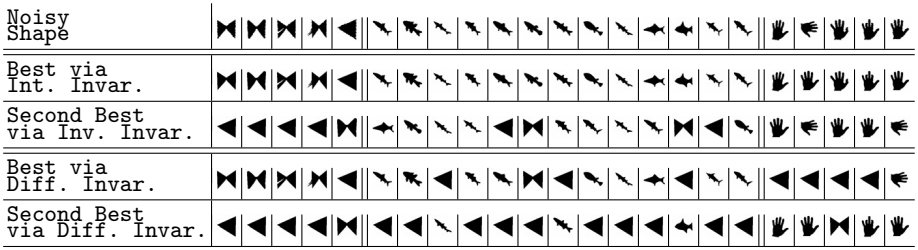

Fig. 7. Summary of noisy shape recognition from a database of 23 shapes.

However, as discussed in Sec. 5 this computation is sensitive to the parameterization of the shapes, specifically the assignment of the initial point. To avoid this dependence, the distance in eq. (9) must be optimized with respect to the choice of $s = 0$. We demonstrate the application of distance computed in this way in the Sec.(7), where we also define a distance based on curvature in the same way.

As an alternative to optimizing D_E, we can define a distance on a parameter-independent representation, such as the signature. The symmetric Hausdorff distance between signature curves (represented as point sets),

$$D_H(\gamma_i, \gamma_j, r) = H((I^{\dot{r}}_{\gamma_i}, I^{r}_{\gamma_i}), (I^{\dot{r}}_{\gamma_j}, I^{r}_{\gamma_j})) \tag{10}$$

is one such distance. Hausdorff distance does not rely on correspondence between points, which is advantageous because it provides the parameter-independent distance we desire, but problematic when non-corresponding segments of the signatures are perturbed so that they overlap.

However, other measures that characterize the signature, such as winding number, can be integrated in into the distance measure to better discriminate these signatures. Additionally, a richer multiscale description of the curve can be created by computing the above distances for a set of kernel sizes. The integration of multiscale information, along with other measures such as winding number, is the subject of ongoing investigation.

7 Experiments

In this section we apply the invariant shape descriptions to the problem of Euclidean-invariant matching of shapes in noise. In Fig. 6, we demonstrate shape matching in a collection of 23 shapes, and summarize the results in Fig. 7. The collection contains several groups of shapes; shapes within a group are similar (i.e. different breeds of fish), but the groups are quite different (intuitively, hands are not like fish).

The figure shows the distance between the shapes (shown on the left side) and noisy versions of the shapes (shown across the top). Within each block are two distances; on top, the integral invariant distance D_E defined in the previous section, and on the bottom the differential invariant distance defined similarly.

In each column, the lowest distance for the shape shown at the top of the column is shown in italics. The distance based on the integral invariant finds the correct match (i.e. the distance between a noisy shape and the correct pair is lowest) in all but one case. The exception is the noisy, rotated hand (fourth column from the right), which has equal distance to itself and its unrotated neighbor, demonstrating the invariance to rotation of this model. Moreover, distances between similar shapes are lower than distances between members of different groups.

Matching results based on the differential invariant are not as consistent as those based on the integral invariant. There are eight mismatches among the 23 noisy images; most frequently, when a shape cannot be matched it is paired with the triangle (fifth from the right). This may be because the curvature of the triangle is zero almost everywhere, and best approximates the mean of many of the noisy curvature functions. More generally, and more problematically, for some groups distances between similar shapes are *higher* than distances between shapes belonging to other groups, violating the required properties of a distance. For instance, the average inter-group distance is 452.8, while the average intra-group distance is 316.6! Compare this to an inter-group distance of 11.0, which is *lower* than the intra-group distance of 17.4 for the integral invariant distance.

8 Conclusion

In this paper we have introduced a general class of integral Euclidean- and similarity-invariant functionals of shape data. We argue that these functionals are less sensitive to noise than differential ones, but can be exploited in similar ways, for instance, to define invariant signature curves that can be used as a representation to define various notions of shape distance. In addition, the integration kernel includes an *intrinsic* scale-space parameter. We presented efficient numerical implementations of these invariants, and, in the limit, established a completeness property for the representation by showing a one-to-one correspondence with curvature. We demonstrated our results with several experiments, including an application to shape matching using synthetic and real data.

References

1. R. Alferez and Y. F. Wang. Geometric and illumination invariants for object recognition. *PAMI*, 21(6):505–536, 1999.
2. K. Arbter, W. E. Snyder, H. Burkhardt, and G. Hirzinger. Applications of affine-invariant fourier descriptors to recognition of 3-d objects. *PAMI*, 12(7):640–646, 1990.
3. S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *PAMI*, 24(4):509–522, 2002.
4. A. Bengtsson and J.-O. Eklundh. Shape representation by multiscale contour approximation. *PAMI*, 13(1):85–93, 1991.
5. M. Boutin. Numerically invariant signature curves. *IJCV*, 40(3):235–248, 2000.
6. R. D. Brandt and F. Lin. Representations that uniquely characterize images modulo translation, rotation and scaling. *PRL*, 17:1001–1015, 1996.

7. A. Bruckstein, N. Katzir, M. Lindenbaum, and M. Porat. Similarity invariant signatures for partially occluded planar shapes. *IJCV*, 7(3):271–285, 1992.
8. A. M. Bruckstein, R. J. Holt, A. N. Netravali, and T. J. Richardson. Invariant signatures for planar shape recognition under partial occlusion. *CVGIP:IU*, 58(1):49–65, 1993.
9. A. M. Bruckstein, E. Rivlin, and I. Weiss. Scale-space semi-local invariants. *IVC*, 15(5):335–344, 1997.
10. E. Calabi, P. Olver, C. Shakiban, A. Tannenbaum, and S. Haker. Differential and numerically invariant signature curves applied to object recognition. *IJCV*, 26:107–135, 1998.
11. D. Chetverikov and Y Khenokh. Matching for shape defect detection. *LNCS*, 1689(2):367–374, 1999.
12. T. Cohignac, C. Lopez, and J. M. Morel. Integral and local affine invariant parameter and applicatioin to shape recognition. *ICPR*, 1:164–168, 1994.
13. J. B. Cole, H. Murase, and S. Naito. A lie group theoretical approach to the invariance problem in feature extraction and object recognition. *PRL*, 12:519–523, 1991.
14. L. E. Dickson. *Algebraic Invariants*. John-Weiley & Sons, 1914.
15. J. Dieudonne and J. Carrell. *Invariant Theory: Old and New*. Academic Press, London, 1970.
16. J. Flusser and T. Suk. Pattern recognition by affine moment invariants. *Pat. Rec.*, 26(1):167–174, 1993.
17. 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. *PAMI*, 13(10):971–991, 1991.
18. D.A. Forsyth, J.L. Mundy, A. Zisserman, and C.M. Brown. Projectively invariant representations using implicit algebraic curves. *IVC*, 9(2):130–136, 1991.
19. L. Van Gool, T. Moons, E. Pauwels, and A. Oosterlinck. Semi-differential invariants. In J. Mundy and A Zisserman, editors, *Geometric Invariance in Computer Vision*, pages 193–214. MIT, Cambridge, 1992.
20. L. Van Gool, T. Moons, and D. Ungureanu. Affine/photometric invariants for planar intensity patterns. *ECCV*, 1:642–651, 1996.
21. J. H. Grace and A. Young. *The Algebra of Invariants*. Cambridge, 1903.
22. C. E. Hann and M. S. Hickman. Projective curvature and integral invariants. *IJCV*, 40(3):235–248, 2000.
23. M. K. Hu. Visual pattern recognition by moment invariants. *IRE Trans. on IT*, 8:179–187, 1961.
24. K. Kanatani. *Group Theoretical Methods in Image Understanding*. Springer, 1990.
25. E. P. Lane. *Projective Differential Geometry of Curves and Surfaces*. University of Chicago Press, 1932.
26. 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. *ICIP*, 1996.
27. G. Lei. Recognition of planar objects in 3-d space from single perspective views using cross ratio. *Robot. and Automat.*, 6(4):432–437, 1990.
28. R. Lenz. *Group Theoretical Methods in Image Processing*, volume 413 of *LNCS*. Springer, 1990.
29. S. Z. Li. Shape matching based on invariants. In O. M. Omidvar (ed.), editor, *Progress in Neural Networks : Shape Recognition*, volume 6, pages 203–228. Intellect, 1999.

30. S. Liao and M. Pawlak. On image analysis by moments. *PAMI*, 18(3):254–266, 1996.
31. T. Miyatake, T Matsuyama, and M. Nagao. Affine transform invariant curve recognition using fourier descriptors. *Inform. Processing Soc. Japan*, 24(1):64–71, 1983.
32. F. Mokhtarian and A. K. Mackworth. A theory of multi-scale, curvature-based shape representation for planar curves. *PAMI*, 14(8):789–805, 1992.
33. D. Mumford, J. Fogarty, and F. C. Kirwan. *Geometric invariant theory.* Springer-Verlag, Berlin ; New York, 3rd edition, 1994.
34. D. Mumford, A. Latto, and J. Shah. The representation of shape. *IEEE Workshop on Comp. Vis.*, pages 183–191, 1984.
35. J. L. Mundy and A. Zisserman, editors. *Geometric Invariance in Computer Vision.* MIT, 1992.
36. L. Nielsen and G. Saprr. Projective area-invariants as an extension of the cross-ratio. *CVGIP*, 54(1):145–159, 1991.
37. P. J. Olver. *Equivalence, Invariants and Symmetry.* Cambridge, 1995.
38. T. Pajdla and L. Van Gool. Matching of 3-d curves using semi-differential invariants. *ICCV*, pages 390–395, 1995.
39. T. H. Reiss. Recognizing planar objects using invariant image features. In *LNCS*, volume 676. Springer, 1993.
40. C. Rothwell, A. Zisserman, D. Forsyth, and J. Mundy. Canonical frames for planar object recognition. *ECCV*, pages 757–772, 1992.
41. C. Rothwell, A. Zisserman, D. Forsyth, and J. Mundy. Planar object recognition using projective shape representation. *IJCV*, 16:57–99, 1995.
42. G. Sapiro and A. Tannenbaum. Affine invariant scale space. *IJCV*, 11(1):25–44, 1993.
43. G. Sapiro and A. Tannenbaum. Area and length preserving geometric invariant scale-spaces. *PAMI*, 17(1):67–72, 1995.
44. J. Sato and R. Cipolla. Affine integral invariants for extracting symmetry axes. *IVC*, 15(8):627–635, 1997.
45. A. Shashua and N. Navab. Relative affine structure: Canonical model for 3d from 2d geometry and applications. *PAMI*, 18(9):873–883, 1996.
46. C. E. Springer. *Geometry and Analysis of Projective Spaces.* Freeman, San Francisco, 1964.
47. Q. M. Tieng and W. W. Boles. Recognition of 2d object contours using the wavelet transform zero-crossing representation. *PAMI*, 19(8):910–916, 1997.
48. J. Verestoy and D. Chetverikov. Shape detect detection in ferrite cores. *Machine Graphics and Vision*, 6(2):225–236, 1997.
49. I. Weiss. Noise resistant invariants of curves. *PAMI*, 15(9):943–948, 1993.
50. A. P. Witkin. Scale-space filtering. *Int. Joint. Conf. AI*, pages 1019–1021, 1983.
51. C. T. Zahn and R. Z. Roskies. Fourier descriptors for plane closed curves. *Trans. Comp.*, 21:269–281, 1972.
52. A. Zisserman, D.A. Forsyth, J. L. Mundy, C. A. Rothwell, and J. S. Liu. 3D object recognition using invariance. *Art. Int.*, 78:239–288, 1995.