ELSEVIER

Image and Vision Computing 15 (1997) 335-344

image and vision COMPUTING

Scale space semi-local invariants

A.M. Bruckstein[a,*], E. Rivlin[a,†], I. Weiss[b,‡]

[a]Department of Computer Science, Israel Institute of Technology, Technion, Haifa, Israel
[b]Computer Vision Laboratory, Center for Automation Research University of Maryland, College Park, MD 20742-3275, USA

Received 3 June 1995; accepted 14 August 1996

Abstract

In this paper we discuss a new approach to invariant signatures for recognizing curves under viewing distortions and partial occlusion. The approach is intended to overcome the ill-posed problem of finding derivatives, on which local invariants usually depend. The basic idea is to use invariant finite differences, with a scale parameter that determines the size of the differencing interval. The scale parameter is allowed to vary so that we obtain a 'scale space'-like invariant representation of the curve, with larger difference intervals corresponding to larger, coarser scales. In this new representation, each traditional local invariant is replaced by a scale-dependent range of invariants. Thus, instead of invariant signature curves we obtain invariant signature surfaces in a 3-D invariant 'scale space'.

Keywords: Local invariants; Object recognition; Scale space

1. Introduction

One of the major problems of object recognition is the fact that, on the one hand, an object can be seen from different points of view, producing different images. On the other hand, we would like to store only one image in a database and match any other image of the object to it, regardless of the point of view. A good way to overcome this problem is to use viewpoint invariants, namely descriptors of the shape that are independent of the point of view, and use them for matching.

The subject of viewpoint invariants in vision has developed rapidly in recent years. A simple projective, or viewpoint, invariant, namely the cross ratio of four points on a line, was introduced in vision by Duda and Hart [1]. However, its domain of applicability was very limited. More general invariants were studied in the nineteenth century, and were introduced in the field of computer vision by Weiss [2]. They are of two main types

(1) *Algebraic invariants.* These are based on a global description of the shapes by algebraic entities such as line, conics and polynomials. Details of these methods can be found in Grace and Young [3] and Springer [4].
(2) *Differential invariants.* These are based on describing the shape by arbitrary differentiable functions. These methods were developed by Halphen [5], Wilczynski [6], Cartan [7] and Lane [8].

These methods have been applied to various vision problems. The algebraic approach was used by Forsyth et al. [9] and Taubin [10], while differential invariants were used by Weiss [2,11] and Bruckstein and Netravali [12]. Both methods proved to have advantages and disadvantages. The algebraic method, while simple and easy to implement, is quite limited in the kinds of shapes that it can handle because most shapes are not representable by simple low order polynomials. The differential method is more general because it can handle arbitrary curves, but it relies on the use of local information such as derivatives (of quite high orders).

This situation has led to the introduction of various kinds of intermediate, or hybrid methods, that try to combine the advantages of the algebraic and differential methods, and hopefully not their disadvantages. Van Gool et al. [13], Brill et al. [14] and others introduced invariants that contain both derivatives and reference points. Each reference point reduces the number of derivatives that one needs in order to obtain invariants.

* Email: freddy@cs.technion.ac.il
† Email: ehudr@cs.technion.ac.il
‡ Email: weiss@cfar.umd.edu

PII S0262-8856(96)01140-7

Weiss [15] used a 'canonical' coordinate system without curve parameterization to obtain the same goal. This resulted in fewer derivatives and in the capability of using feature lines in addition to points. However, in all these methods, the correspondence must be established between the reference points of the two images that are being matched. Finding the correspondence is a very difficult problem that requires searches in high dimensional spaces, and we need a method that avoids this.

In this paper we reduce the number of derivatives by using a scale space approach. It is well known [16,17] that such an approach can turn the ill-posed problem of finding derivatives into a well-posed one. The scale space has to be invariant, so we cannot use simple Gaussian-like smoothing. Instead, we rely on some reference points as a function of the given curve and a variable scale parameter. These reference points are not assumed to be readily available in the image, as in previous methods [18,19,21], but are determined from the curve in an invariant way. Thus, no correspondence is needed. Using low-order derivatives and our variable reference points, we build invariant scale space representations of the given curves.

There are various ways to derive invariants in accordance with the above scheme. Here we extend a method originally introduced by Bruckstein et al. [20,22]. It consists of defining an invariant arclength (using the lowest possible order of derivatives in given schemes), and then defining invariant finite differences using this arclength. These differences replace the higher order derivative in the traditional invariants. The differences are not necessarily small and do not tend to zero. Rather, their variable size creates the 'scale space'.

We briefly describe here an illustrative example of the method. Given a curve, we want to find invariants at each point of the curve so that we can obtain a local invariant signature. With Euclidean invariance in mind, we can plot the curvature vs. the arclength τ to obtain a Euclidean invariant signature. Invariant signature plots of two curves are then compared to detect matches rather than the curves themselves. This is an example of a local method, in which no correspondence between points is needed. However, curvature involves a second derivative which we wish to avoid.

In our new method, the second derivative is replaced by a finite difference. We start from a point $P(\tau)$ on the curve, and we want to find invariants there. We choose an interval size $\Delta\tau$ and find two points on the curve, $P(\tau+\Delta\tau)$, $P(\tau-\Delta\tau)$, located at distances $+\Delta\tau$ and $-\Delta\tau$ (measured on the curve) from the point $P(\tau)$ at which we want to calculate the invariants. Given these three points, we can calculate any Euclidean invariant involving them, such as the area $A(\tau)$ of the triangle formed by them. $A(\tau)$ is then a new type of invariant signature. This is much more robust than a derivative, if $\Delta\tau$ is not too small. In this way, we reduce the number of derivatives needed without needing any fixed reference points or their correspondence. The scale parameter $\Delta\tau$ can now be varied to obtain a whole range of scale dependent invariants.

Similar difference-based methods have been used earlier in various contexts. Euclidean variants have been used for detecting inflection points and other features on boundaries [23]. A surface description using a 'tripod', or a triangle with a known side h which is superimposed on the surface, was used by Pipitone [24], and proved to be quite robust.

In summary, the semi-local, or finite difference method elaborated upon in Bruckstein, Holt et al. [22] is extended here as follows. As in [22], we consider general transformations such as similarity, affine, or even projective viewing distortion and use similarity, affine or projective invariant arclength to reparametrize the curve, exploiting all the information available. *We then let the differencing interval size or sizes be free parameters rather than setting them in advance.* In this way we obtain whole ranges of invariants at each point rather than single values. The signature functions for the curves then become signature vectors or even continua of values, i.e. surfaces or hypersurfaces. Matching them will be slightly more complicated but will certainly be robust because it will be less sensitive to peculiarities that may exist at some fixed pre-set value of the locality (scale) parameters.

2. Theory of scale dependent local invariants

Here we describe in detail the basic ideas of the semi-local method. Its main advantage over the global method is its ability to deal with partially occluded shape. We deal here with planar curves, such as boundaries of planar objects.

To obtain an invariant representation of a curve, we associate with each point of the curve a set of invariants. The collection of independent invariants from all points is the invariant 'signature'. This approach maps the problem into a problem of detecting partial matches between the signatures of the 'library' of possible objects and the signature functions extracted from the (composite) objects appearing in the scene to be analysed.

We treat here two variants of such signatures:

- A signature with an arclength as an independent variable. We first derive an invariant arclength τ and use it to reparametrize the curve. After that another invariant I is determined at each point (for example, curvature in the Euclidean case), and we represent the signature as a function $I(\tau)$.
- A signature with two independent absolute invariants as coordinates. Here we find two local invariants I_1, I_2

at each curve point, and then plot I_2 against I_1. The functions I_1, I_2 may or may not be represented as functions of an invariant arclength τ, but they are local quantities in any case.

In both cases, recognition is based on detecting portions of invariant traces in the 'transform plane'.

The simple Euclidean example will again be used to clarify the above discussion. Suppose we wish to detect the presence of partially occluded planar objects whose instances may undergo planar rotations and translations (i.e. Euclidean transformations). Here the well-known invariant signature approach describes object boundaries via curvature versus arclength functions, invariant under Euclidean transformation, and recognition is possible by partial matching. This method of finding a signature is based on using an invariant metric on the curve (the arclength) and on finding a differential invariant at each point on the curve (the curvature).

The second approach above was used in Refs. [15,31], without a scale space. No curve parametrization was used there. The method is based on the fact [29] that a general curve has two independent invariants at each point, and these can determine the curve uniquely up to the relevant transformation. Here we describe various ways of using the method in scale space. We deal with the case of Euclidean invariance. One way to proceed is to use a Euclidean arclength τ. Associate with points on the arbitrarily parametrized curve $P(\tau) = [x(r), y(\tau)]$ two numbers $I_1(\tau)$ and $I_2(\tau)$ invariant under Euclidean transformation. Here, $I_1(\tau)$ could be the curvature at $P(\tau)$ and $I_2(\tau)$ could be the area of the triangle formed by the points $\{P(\tau), P(\tau - \tau_b), P(\tau + \tau_f)\}$, where $P(\tau - \tau_b)$ is located at an *arclength* distance of τ_b (chosen a priori) 'before' $P(\tau)$ and $P(\tau + \tau_f)$ at a distance of τ_f 'following' $P(\tau)$ in the traversal of the curve (see Fig. 1).

A more appealing way to use the second method is to avoid the curve parameter altogether. We can avoid it by a variety of methods. For example, we can define the first invariant at $P(\tau)$ to be the area between the curve and a parallel to the tangent at $P(\tau)$ at a distance of D (set beforehand) toward the center of the osculating circle $C(\tau)$ (see Fig. 2).

To summarize our message: the first, metric-based approach to finding a signature, calls for the arclength reparametrization of the curve, i.e. $P(t) \rightarrow P(\tau)$, and the

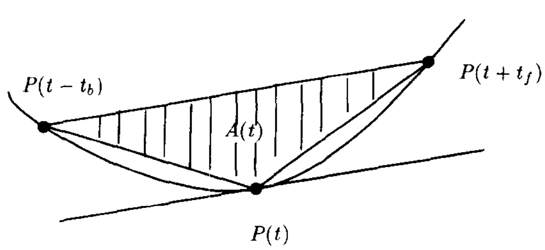

Fig. 1. $P(t) \rightarrow [I_1(t) = 1/R(t), I_2(t) = A(t)]$.

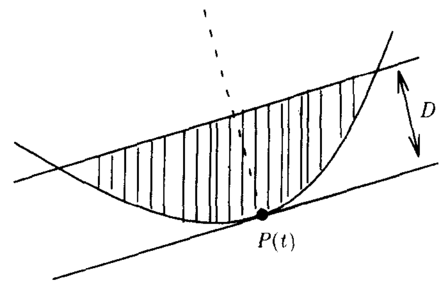

Fig. 2. Using the area between the curve and a parallel to the tangent.

association of one invariant quantity with each point of the reparametrized curve (in the above example — the curvature). The second, invariant coordinate approach, associates different invariant quantities with each point of the curve without necessarily referring to a curve parametrization.

Note that both approaches are based on our ability to analyse the neighborhood of a point on a planar curve and calculate some quantity that remains the same when we consider the *image of the point* and *the image of its neighborhood* under the viewing transformation.

In this paper we concentrate on the first method discussed above, i.e. we assume that we can always determine an invariant metric on the curve. With this metric, moving to the left and right along the curve from a point $P(\tau)$ to points at 'distances' $\pm\Delta\tau_1, \pm\Delta\tau_2, \pm\Delta\tau_3, \ldots$, etc., is a well-defined process. This process can be used to generate point sets anchored at $P(\tau)$ that are invariant under the distorting viewing transformation. Based on these point sets, we are able to use the global invariants of the viewing transformation to calculate a wide variety of invariants.

More importantly, notice that the point sets are parametrized by the sequence of positive numbers $\Delta\tau_1 < \Delta\tau_2 < \Delta\tau_2 < \cdots$ and hence the invariant quantities that we generate are likewise parametrized. We can use this freedom to associate with each point $P(\tau)$ on the curve a whole range of invariants rather than a single one. Hence we can define *multi-valued* or *parametrized* signatures (or coordinates in the second approach discussed above). These have the potential of enabling more robust matching in the presence of noise and other disturbances.

To illustrate this let us again consider the Euclidean case. Once the curve $P(t)$ is reparametrized to $P(\tau)$ where τ is Euclidean arclength we can proceed as follows: at each point $P(\tau)$, consider the point set $\{P(\tau - \Delta\tau), P(\tau), P(\tau + \Delta\tau)\}$ and compute the radius of the circle passing through these points, denoting it by $R(\tau, \Delta\tau)$. Clearly, as $\Delta\tau \rightarrow 0, R(\tau, \Delta\tau) \rightarrow 1/k(\tau)$ where $k(\tau)$ is the curvature. However, we can use the whole range of values of $\Delta\tau$ from $\Delta\tau = 0$ to some $(\Delta\tau)_{max}$ to associate with $P(\tau)$ a multi-valued curvature function of the form $k(\tau, \Delta\tau) = 1/R(\tau, \Delta\tau)$. At $\Delta\tau = 0$ we obtain $k(\tau)$, but $k(\tau, \Delta\tau)$ clearly carries more information on the local behavior of the curves around $P(\tau)$ than $k(\tau)$, in the neighborhood of any value of τ.

Furthermore, we can use other Euclidean invariant quantities, like the areas of the triangles $\{P(\tau), P(\tau-\Delta\tau), P(\tau+\Delta\tau)\}$, i.e. $A(\tau,\Delta\tau)=$ Area $\{P(\tau), P(\tau-\Delta\tau), P(\tau+\Delta\tau)\}$, the angles $\varphi(\tau,\Delta\tau)=\angle P(\tau-\Delta\tau)P(\tau)P(\tau+\Delta\tau)$, etc. All these are valid 'generalized parametrized signature' functions that can be associated with a planar curve. (There are clearly relationships between the various quantities, but this will not concern us here.) In case we need to recognize occluded planar shapes under a Euclidean viewing transformation these generalized signatures will enable us to perform more robust partial matching for detection.

Note that the same approach can also be used in conjunction with the invariant coordinate method. One of the invariant quantities can be chosen as the independent variable $\tau = I_1(t)$, and the other can be a parametrized continuum of values $I_2(\tau,\Delta\tau)$.

Note also the important point that we do *not* necessarily advocate the computation of the limit values for $\Delta\tau \to 0$. If $\Delta\tau$ takes only a finite set of positive values we base our invariants on a form of *finite differences* in the invariant metric, rather than on the differential behavior of the curve about $P(\tau)$.

3. Scale space of invariants under similarity, affine and projective maps

So far we have illustrated our approach with examples based on the simplest case of Euclidean transformations affecting planar objects. If more complex viewing transformations are assumed we must deal with such problems using transformation-invariant metrics on curves and using the arsenal of geometric invariants available for the specific transformations.

Much mathematical research was devoted to the famous Klein program in the 19th century (e.g. see Buchin [25], Lane [26], Wilczinsky [6,27,28] and Guggenheimer [29]), who show how to compute invariant reparametrizations under various continuous groups of transformations. The affine and projective groups have received special attention in this context, and we have almost 'off-the-shelf' results available from affine and projective differential geometry.

Given a planar curve $P(t)$, with arbitrary parametrization, the local behavior of this curve at a point $P(t_0)$ is described by the vector of derivatives $\{P(t), P'(t), P''(t)\ldots P^{(n)}(t)\ldots\}_{t=t_0}$. Suppose now that we are given a transformed image of $P(t)$ — arbitrarily parametrized by $\tilde{t}$, i.e. $\tilde{P}(\tilde{t}) = T_\psi\{P(\tilde{t}(t))\}$. The local behavior of $\tilde{P}(\tilde{t})$ at the image of $P(t_0)$, i.e. $\tilde{P}(\tilde{t}(t_0)) = \tilde{P}(\tilde{t}_0)$, will clearly be related to the local behavior of $P(t)$ at t_0. However, the relationship is quite difficult to describe analytically due to the arbitrariness of the parametrizations and the complexity of the (possibly nonlinear) viewing transformation T_ψ.

Suppose that, in spite of such difficulties, we can produce some functionals $\Omega[\cdot]$ associating numbers with points on the curve such that

$$\Omega[P(t_0)] = \Omega[\tilde{P}(\tilde{t}_0)] \tag{1}$$

i.e. the numbers are invariant under the transformation $T_\psi\{\cdot\}$ and the necessarily arbitrary parametrizations. If two such functions $\Omega_1[]$ and $\Omega_2[]$ are available and are not trivially related (for example we don't have $\Omega_1[] = f\{\Omega_2[]\}$), then we can base a recognition procedure on them via an invariant coordinate method.

Indeed, we can simply associate

$$P(t) \to [I_1(t) = \Omega_1[P(t)], I_2(t) = \Omega_2[P(t)]]$$

and the traversal of $P(t)$ yields an invariant trace in the $[I_1, I_2]$ plane.

Furthermore, by differentiation of (1) $\Omega[P(t)] = \Omega[\tilde{P}(\tilde{t}(t))]$ w.r.t. to t we obtain

$$\frac{\mathrm{d}}{\mathrm{d}t}\Omega[x(t)] = \frac{\mathrm{d}}{\mathrm{d}\tilde{t}}\Omega[\tilde{P}(\tilde{t})]\cdot\frac{\mathrm{d}\tilde{t}}{\mathrm{d}t}$$

and defining $\Gamma[P(t)] \triangleq \mathrm{d}/\mathrm{d}t\,\Omega[P(t)]$ we can write

$$\Gamma[P(t)]\,\mathrm{d}t = \Gamma[\tilde{P}(\tilde{t})]\,\mathrm{d}\tilde{t} \tag{2}$$

This enables us to reparametrize the curves $P(t)$ and $P(\tilde{t})$ using an invariant metric, since defining $d\tau \triangleq |\Gamma[P(t)]|dt$ yields

$$d\tilde{\tau} = |\Gamma[\tilde{P}(\tilde{t})]|d\tilde{\tau} = |\Gamma[P(t)]|dt = d\tau$$

This is just like in the formula for arclength reparametrization

$$d\tilde{\tau} = \sqrt{(\tilde{x}^1(\tilde{t}))^2 + (\tilde{y}^1(\tilde{t}))^2}\,d\tilde{t}$$
$$= \sqrt{(x^1(t))^2 + (y^1(t))^2}\,dt = d\tau$$

for the Euclidean group of transformations.

Note that τ, like the Euclidean arclength, is a 'monotonized' version of Ω, and that it depends on an arbitrary parameter: the initial value, or the starting point of the integration along the planar curve.

Thus, the question is: how can we obtain nontrivial functions $\Omega[\cdot]$ (or $\Gamma[\cdot]$) obeying Eqs. (1) or (2)? Note that these functions should reflect the geometry of the curve around a point of interest, and here 'around' means an infinitesimal neighborhood (i.e. the differential geometry at the point of interest.)

In Bruckstein and Netravali [12], Weiss [30], Bruckstein, Katzir et al. [20] and Bruckstein, Holt et al. [22], several such functions are exhibited for all the viewing transformations. It is also pointed out there that the requirement for $\Omega[]$ or $\Gamma[]$ to be based on infinitesimal behavior can be relaxed, and that the use of non-local properties of the viewing transformation

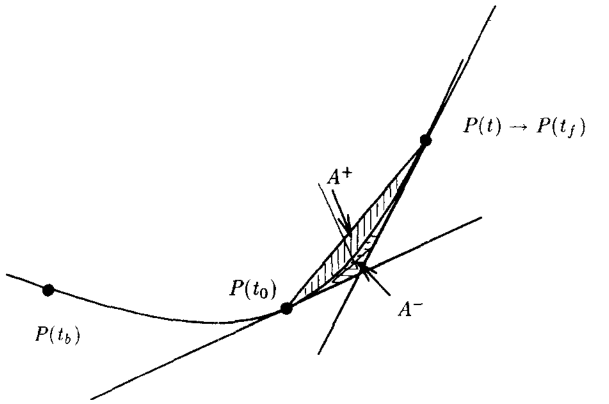

Fig. 3. An invariant parametrization using the area ratios.

could be used to define such functions over 'invariantly defined' finite neighborhoods of points on planar curves. As an example, let us consider curves distorted by affine transformations.

First we find an invariant parametrization of the curve. This can be done by several methods. In one method, depicted in Fig. 3, we can consider the area ratios $A^+(P(t))/A^-(P(t))$ for $P(t)$ around $P(t_0)$, i.e. for $t > t_0$ and $t < t_0$. In Bruckstein, Holt et al. [22], it was proposed to look for $P(t)$ at $t > t_0$ such that A^+/A^- first equals some constant k_F and for $P(t)$ at $t < t_0$ such that A^+/A^- first equals some constant k_B. Then the ratio of the corresponding areas is an affine invariant quantity dependent on the two parameters k_F and k_B,

$$\Omega_1(x(t_0)) = \frac{A_F(k_F)}{A_B(k_B)}$$

(We denote $A = A^+ + A^-$.) As a generalization of this work, we can let $k_F = k_B$ take a whole *range* of values, associating a range of invariants with the point $P(t_0)$. We can also consider another invariant Ω_2, the ratio of the area of $(P(t_0)P(t_F)P(t_B))$ to the area $A_F + A_B$.

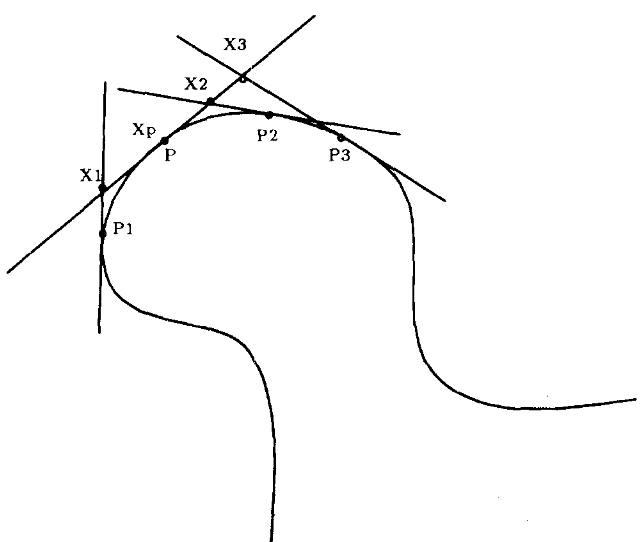

Fig. 4. Obtaining projective invariants.

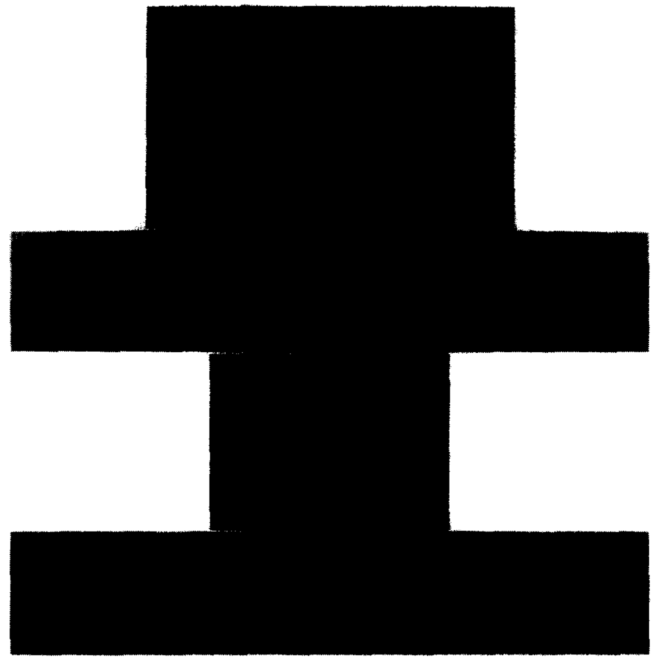

Fig. 5. A conic transformed by scaling and rotation. The multi-valued signature for each of the conics is presented below it. The x-axis represents position along the reparametrized curve. Signatures for 20 different parameter values are displayed along the y-axis. The starting position for each curve is marked by a white square. A match is achieved when one of the signatures is shifted.

Another method for affine invariant parametrization is to use a differential formula, after appropriate smoothing of the curve. We have the invariant arclength [29]

$$d\tau = |(x_t y_{tt} - x_{tt} y_t)^{1/3}| dt$$

with the subscripts denoting derivatives with respect to t. We perform the smoothing by a spline interpolation over some interval Δt. From this, using the above formula, it is easy to calculate the corresponding invariant parameter $\Delta\tau$.

For each $\Delta\tau$ as calculated in any of the methods above, we have a set of points $P(\tau)$, $P(\tau + \Delta\tau)$, $P(\tau - \Delta\tau)$. For such a set we can define affine invariants as ratios of areas defined by these points, the tangents and the curve itself. Since $\Delta\tau$ is a free scale parameter, we obtain a whole range of invariants associated with the point $P(\tau)$.

We turn to the projective case. Here too we can find a projective arclength using a differential formula, again with appropriately smoothing with splines. This is done following Wilczynski [6]. A projective transformation

Fig. 6. Three logos from the logo database.

Fig. 7. A logo before and after transformation (scaling and rotation). The logo was processed as five different curves.

can be written in homogeneous coordinates as

$$\tilde{\mathbf{x}} = \lambda(\mathbf{x})T\mathbf{x}$$

with $\lambda(\mathbf{x})$ being an arbitrary factor, which can be different at each point $\mathbf{x}$. To find invariants, one can proceed in stages. First find invariants to the linear part T of the transformation, and from those derive invariants to λ, and also to change in the curve parameter t.

Given a plane curve $\mathbf{x}(t)$, invariants to T can be obtained by solving the linear algebraic system of equations

$$\mathbf{x}''' + 3p_1\mathbf{x}'' + 3p_2\mathbf{x}' + p_3\mathbf{x} = 0$$

for the three unknowns, p_1, p_2, p_3, at each point t, where the primes indicate derivatives with respect to t. It is easy to show, by multiplying the equation through by T, that these solutions p_i are invariant to T. However, they are not invariant to change in the arbitrary factor $\lambda(\mathbf{x}(t))$ nor to change in the curve parameter t. We can

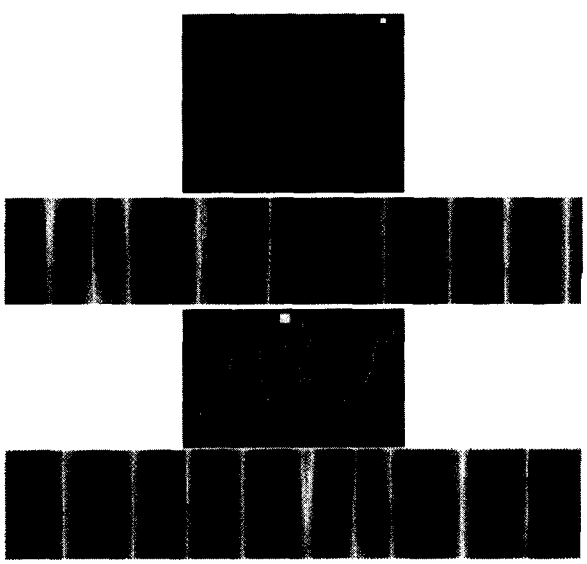

Fig. 8. The last two letters of the logo (*gs*) constitute a curve to be processed. The multi-valued signatures are presented below. Shifting the top signature to the right by approximately half the strip's length will achieve a match.

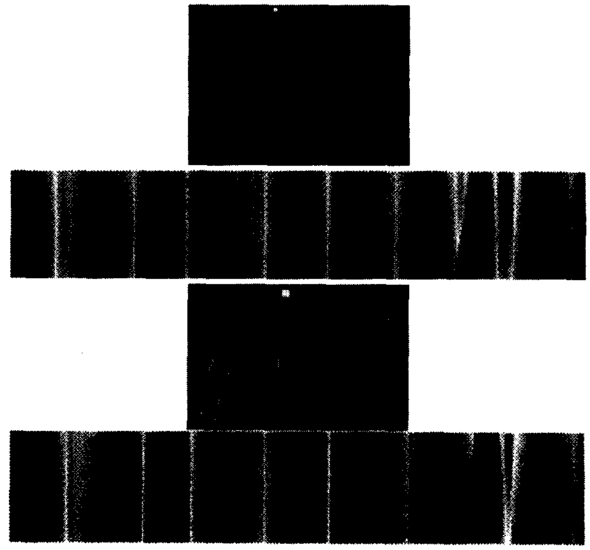

Fig. 9. A second curve from the logo, taken from the first four letters *Kell*). Again, a match is achieved after a shift.

obtain functions of these p_i which are invariant to the additional possible transformation. We have the 'semi-invariants'

$$\bar{p}_2 = p_2 - p_1^2 - p_1'$$

$$\bar{p}_3 = p_3 - 3p_1p_2 + 2p_1^3 - p_1''$$

These remain unchanged under multiplication of the coordinates by a factor $\lambda(\mathbf{x})$, but not under change of the parameter t.

An invariant arclength can be defined as $d\tau = |(\bar{p}_3 - \frac{3}{2}\bar{p}_2')^{1/3}|dt$. This is an absolute invariant with respect to changing the parameter as well as the projection, except for a starting point. We can now reparametrize the curve to obtain $P(\tau)$. Then, we generate a range of invariants around each $P(\tau)$, varying according to a scale parameter $\Delta\tau$. An example is shown in Fig. 4. We draw points on the curve separated by intervals $\Delta\tau$, and draw the tangents at these points. The tangents produce a set of points on a straight line from which we calculate the cross ratio:

$$CR = \frac{(x_2 - x_1)(x_P - x_3)}{(x_P - x_1)(x_2 - x_3)}$$

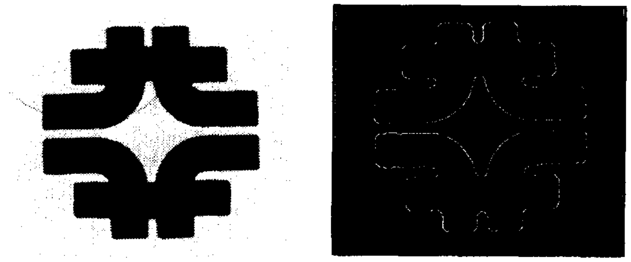

Fig. 10. A second input logo. Only the lower curve is processed.

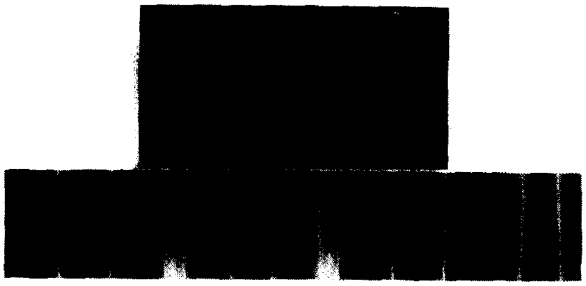

Fig. 11. The lower curve and its signature. One can see the symmetry of the shape from the structure of the signature.

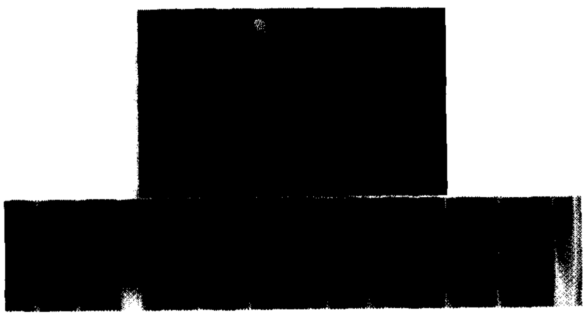

Fig. 12. The lower curve after scaling and rotation.

This invariant depends on the point $P(\tau)$, and also on the scale parameter $\Delta\tau$.

4. Experiments: invariant scale space signatures

We present a series of experiments to illustrate the above-outlined theory. We start by considering a simple conic, and we try to obtain its scale-space invariants under the similarity group of transformations. The conic before and after it went through scaling and rotation is presented in Fig. 5.

To obtain an invariant arclength under similarity we proceed as in Bruckstein, Katzir et al. [20]. The similarity invariant arclength parameter is given in this case by

$$d\tau = \left| \frac{x'(t)y''(t) - x''(t)y'(t)}{(x'(t))^2 + (y'(t))^2} \right| dt$$

After the curve is reparametrized by the invariant arclength we can call upon several types of scale-dependent similarity invariants. In this example (and in the ones that follow), we plot the angle $[P(\tau - \Delta\tau)P(\tau)P(\tau + \Delta\tau)] = \varphi(\tau, \Delta\tau)$ as a function of τ. However, a wealth of other possibilities are available. We could also compute various length or area ratio that are also known to be similarity invariants.

The multi-valued signature for each of the conics is presented in Fig. 5. The invariant arc length is represented by the x-axis, which represents position along the reparametrized curve. The y-axis represents the values of the scale parameters $\Delta\tau$. In our experiments each image contains 20 different signatures for 20 different parameter values. For each signature different $\Delta\tau$s were used. For a constant Y value one gets single-valued signatures for the curve. The grey level encodes the similarity invariant for a particular arclength and parameter value. The full display represents an 'invariant signature surface'. For each curve the starting position is marked by a white square. Due to the different starting position

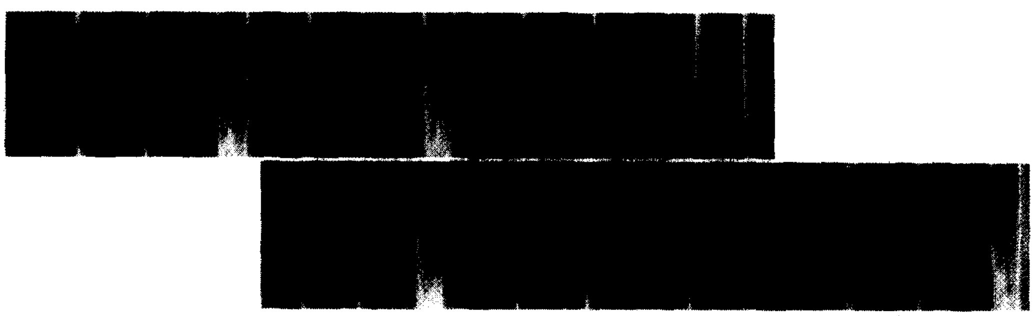

Fig. 13. Shifting the lower curve's multi-valued signature to the right gives the highest correlation value, and achieves a match.

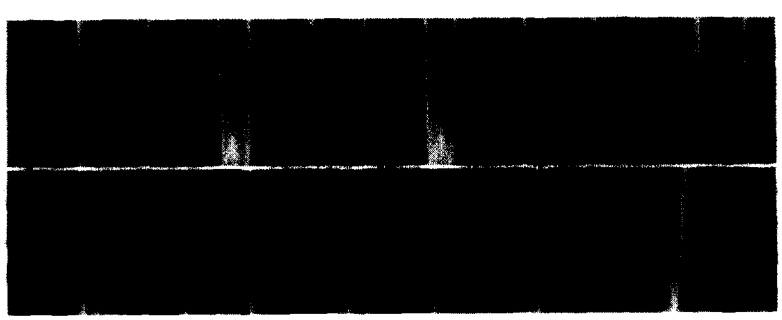

Fig. 14. The multi-valued signature from Fig. 11 (top) is compared with the multi-valued signature from Fig. 9. No match is achieved.

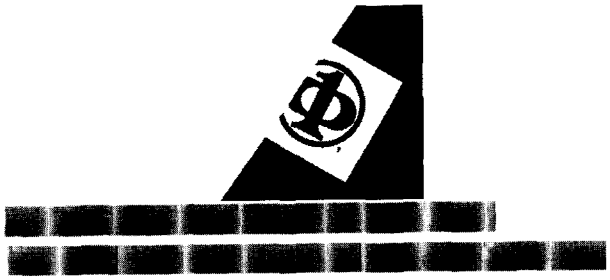

Fig. 15. An occluded logo and its multi-valued signature. The signature of the complete logo from the database is presented below. The presented relative position of the two signatures is the one which gives the best match.

one multi-valued signature is shifted relative to the other. To check for a match between two signatures one should match one multi-valued signature to the other while shifting it in a cyclic manner. This is done automatically, scoring each match, trying to solve for the maximum. In this example, one can see that a match is achieved when one of the signatures is shifted. The symmetric nature of the curve is evident in the signature structure.

For the following experiments we used images of different logos. We regard logo recognition as a good application, where similarity invariants are sufficient for recognition. We used the logo database of the Document Processing Group at the University of Maryland. Examples of the input images for the process are presented in Fig. 6. Each logo is processed to an edge image. Curves are processed by length, and a B-spline is matched to each of the curves. We use a matching technique similar to the one presented in Pauwels et al. [32]. Using the B-spline, derivatives are computed, and invariant arclength is obtained.

In Figs. 7–14 we show the results of the processing on two images of logos going under similarity transformation. In Fig. 7, a familiar logo was processed and mapped into five different curves. The scale-space signatures for two curves out of the five are shown in Figs. 8 and 9. Each position on the base signature was checked

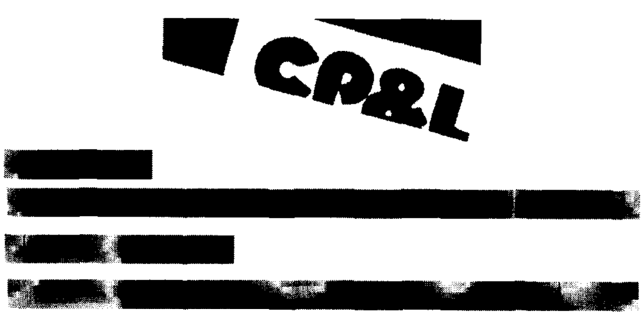

Fig. 16. The logo was processed as four different curves. On top the multi-valued signature for the letter P under occlusion, and below it the signature for the complete curve of the letter P. The signatures below belong to the sign &. The third signature from the top for the occluded curve, and the complete curve below it. The presented relative position of the two signatures is the one which gives the best match.

Fig. 17. An input logo under affine transformation.

against the compared signature in a circular process. Difference values for each position were calculated. In both cases a good match is achieved modulo a (circular) shift in the invariant arclength. The signatures as presented in the figures are showed in the position that gave the best value for the match (minimum of differences between the signatures). The computation time on a Sparc IPX station took an average of three minutes per logo.

The results for a second logo (Fig. 10) are presented in Figs. 11 and 12. The shape is symmetric and only one of the two curves comprising it (the lower) was processed. The processed curve was itself symmetric. One can see the symmetry in the shape from the structure of the signature. Again a good match is achieved when one of the signatures is shifted.

In the following, we show results of experiments that handle occlusion. The processing stage is similar. When we produce signatures for open curves using different parameters we have different domain for each parameter. Hence, extracting multi-valued signatures from occluded curves forces us to further reduce the common domain. As a result we are restricted with the amount of occlusion we can handle without difficulty. Still occlusion of 30% of the shape can be handled without any problem. The automatic matching is done in the same manner as before, comparing the multi-valued signature in each step in a simple cyclic move. The best score achieved is the result of the comparison. In this experiment each multi-valued signature contains only five different signatures for five different parameter values.

Figs. 15 and 16 show two logos under occlusion. The multi-valued signatures of the occluded curves are presented below the images of the logos. One can see that a good match is obtained for the two logos.

We shall next deal with affine transformations. In this case, as noted before, the invariant arclength parameter can be obtained from any arbitrary parameter s by

$$d\tau = |(x_t y_{tt} - x_{tt} y_t)^{1/3}|\,dt$$

with the subscripts denoting derivatives with respect to t.

Fig. 18. The multi-valued signature for the letter F. On top the signature for the database logo, below the signature for the logo after the transformation. The presented relative position of the two signatures is the one which gives the best match.

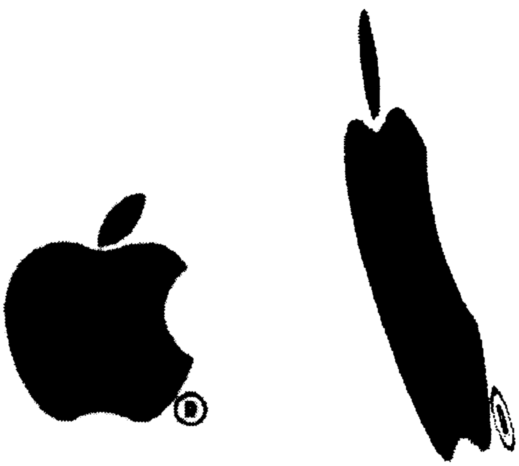

Fig. 19. An input logo under affine transformation.

This expression is an equi-affine invariant. If a full affine transformation is applied, this expression is invariant up to a factor equal to the determinant of the transformation. To make the arclength fully invariant in this case too, we could normalize the expression above by the total arclength $\int d\tau$ of the curve segment we deal with, provided it is not partially occluded. Alternatively, we could use higher derivatives or point matching methods to obtain similar types of invariant arclength.

After the curve is reparametrized by the invariant arclength we can call upon several types of affine invariants. In the examples shown (Figs. 17–20), we plot areas ratio similar to the one presented in Section 3 against the invariant arclength parameter. In this experiment too, each multi-valued signature contains five different signatures for five different parameter values.

5. Discussion

We have developed a way of improving the reliability of object recognition by the method of local invariants. The advantage of local invariants relative to global ones is their ability to handle occlusion. The difficulty in using them lies in the need to use derivatives. Derivatives are not very robust to noise, and even in the noiseless case they can depend upon the scale at which we look at the image, namely the degree of smoothing. We have proposed solving this problem by looking at the shape at many scales rather than trying to choose one particular scale factor, which is not invariant.

Instead of derivatives, we use a finite difference method in an invariant form. The differences depend on a scale factor which we allow to vary continuously, thereby obtaining a description of the shape in an invariant scale space. Scale space methods have been extensively used, but mostly not in an invariant way. The treatment here is quite general; several forms of difference-based invariants have been treated here for

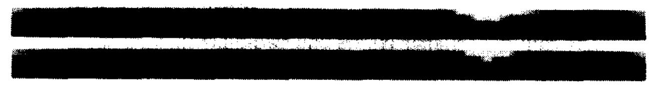

Fig. 20. The multi-valed signature for the Apple logo. On top the signature for the database logo, below the signature for the logo after the transformation. The presented relative position of the two signatures is the one which gives the best match.

projective, affine and similarity transformations. We have shown experimentally that the method can easily recognize various complicated shapes.

Acknowledgements

The support of the Air Force Office of Scientific Research under Grant F49620-92-J-0332 is gratefully acknowledged, as is the help of Sandy German in preparing this paper.

References

[1] R.O. Duda and P.E. Hart, Pattern Classification and Scene Analysis, Wiley, New York, 1973.
[2] I. Weiss, Projective invariants of shapes, Proc. DARPA Image Understanding Workshop, Cambridge, MA, 1988, 1125–1134.
[3] J.H. Grace and A. Young, The Algebra of Invariants, Chelsea, New York, 1903.
[4] C.E. Springer, Geometry and Analysis of Projective Spaces, Freeman, San Francisco, 1964.
[5] M. Halphen, Sur les Invariants Différentiels des Courbes Gauches, J. Ecole Polytechnique, 28(1) (1880).
[6] E.J. Wilczynski, Projective Differential Geometry of Curves and Ruled Surfaces, Teubner, Leipzig, 1906.
[7] E. Cartan, La Théorie des Groupes Continus et la Géometrie, Oeuvres Complètes, III/2, Gauthier-Villars, Paris, 1995, 1727–1861.
[8] E.P. Lane, A Treatise on Projective Differential Geometry, University of Chicago Press, 1942.
[9] D. Forsyth, J.L. Mundy, A. Zisserman and C.M. Brown, Projectively invariant representations using implicit algebraic curves, Image and Vision Computing, 8 (1990) 130–136.
[10] G. Taubin and D.B. Cooper, Object recognition based on moment (or algebraic) invariants. In: Geometric Invariance in Machine Vision, J.L. Mundy and A. Zisserman (eds.), MIT Press, Cambridge, MA, 1992, 375–397.
[11] I. Weiss, Geometric invariants and object recognition, Int. J. Computer Vision (accepted).
[12] A.M. Bruckstein and A.N. Netravali, On differential invariants of planar curves and the recognition of partially occluded planar shapes, AT&T Technical Memo, July 1990.
[13] L. Van Gool, P. Kempenaers and A. Oosterlinck, Recognition and semi-differential invariants, Proc. CVPR, 1991, 454–460.
[14] E. Barrett, P. Payton, N. Haag and M. Brill, General methods for determining projective invariants in imagery, CVGIP:IU, 53 (1991) 45–65.
[15] I. Weiss, Local Projective and Affine Invariants, TR 612, Center for Automation Research, University of Maryland, 1992.
[16] J.J. Koenderink and A.J. Van Doorn, Affine invariants from motion, J. Opt. Soc. Am. A (1991) 377–385.
[17] A.P. Witkin, Scale-space filtering, Proc. IJCAI, 1983, 1019–1022.
[18] H. Asada and M. Brady, The curvature primal sketch, IEEE Trans. PAMI, 8 (1986) 2–14.

[19] F. Mokhtarian and A. Mackworth, Scale-based description and recognition of planar curves and two-dimensional shapes. IEEE Trans. PAMI, 8 (1986) 34–43.
[20] A.M. Bruckstein, A. Katzir, M. Lindenbaum and M. Porat, Similarity Invariant Recognition of Partially Occluded Planar Curves and Shapes. Technion CIS Report, June 1990.
[21] K. Kanantani, Tracing planar surface motion from a projection without knowing the correspondence. CVGIP, 29 (1985) 1–12. (CVGIP:IU, 53 (1991) 45–65.
[22] A. Bruckstein, J. Holt, A.N. Netravali and T.J. Richardson, Invariant signatures for planar shape recognition under partial occlusion. CVGIP:IU, 58 (1993) 49–65.
[23] J.H. Han and T. Poston, Distance accumulation and planar curvature. Proc. ICCV (1993) 487–491.
[24] Pipitone, Personal communication.
[25] S. Buchin, Affine Differential Geometry, Science Press, Beijing, China and Gordon & Breach Science, New York, 1983.
[26] E.P. Lane, Projective Differential Geometry of Curves and Surfaces, University of Chicago Press, 1932.
[27] E.J. Wilczynski, Projective differential geometry of curved surfaces (first memoir). Amer. Math. Soc. Trans., 8 (1907).
[28] E.J. Wilczynski, Projective differential geometry of curve surfaces (second memoir). Amer. Math. Soc. Trans., 9 (1908).
[29] H. Guggenheimer, Differential Geometry, Dover, New York, 1963.
[30] I. Weiss, Noise resistant invariant of curves. In: Geometric Invariance in Machine Vision, J.L. Mundy and A. Zisserman (eds.), MIT Press, Cambridge, MA, 1992, 135–156.
[31] E. Rivlin and I. Weiss, Local invariants for recognition. IEEE Trans. PAMI, 17 (1995) 226–238.
[32] E.J. Pauwels, T. Moons, L.J. Van Gool and A. Oosterlinck. Recognition of planar shapes under affine distortion. Int. J. Computer Vision, 14 (1995) 49–65.