



ELSEVIER

Pattern Recognition Letters 20 (1999) 635–650

Pattern Recognition
Letters

www.elsevier.nl/locate/patrec

Corner detection via topographic analysis of vector-potential

Bin Luo^{a,b}, A.D.J. Cross^a, E.R. Hancock^{a,*}

^a *Department of Computer Science, University of York, York YO1 5DD, UK*

^b *Anhui University, Anhui, People's Republic of China*

Received 22 September 1998; received in revised form 23 December 1998

Abstract

This paper describes how corner detection can be realised using a new feature representation based on a magnetostatic analogy. The idea is to compute a vector-potential by appealing to an analogy in which the Canny edge-map is regarded as an elementary current density residing on the image plane. In this paper, we demonstrate that corners are located at the saddle-points of the magnitude of the vector-potential. These points correspond to the intersections of saddle-ridge and saddle-valley structures, i.e. to junctions of the edge and symmetry lines. We describe a template-based method for locating the saddle-points. This involves performing a non-minimum suppression test in the direction of the vector-potential and a non-maximum suppression test in the orthogonal direction. Experimental results using both synthetic and real images are given. We investigate the angle and scale sensitivity of the new corner detector and compare it with a number of alternative corner detectors. © 1999 Elsevier Science B.V. All rights reserved.

Keywords: Corner detection; Topographic analysis; Vector-potential; Saddle-point detection

1. Introduction

Corners are important dominant points in digital images. In many computer vision tasks, such as image registration, image matching (Costabile et al., 1985), object recognition (Liu and Srinath, 1990b; Han and Jang, 1990) and motion analysis (Dreschler and Nigel, 1992), accurate corner detection is essential. Broadly speaking, there are two corner detection strategies adopted in the literature. The first of these is based on the analysis of pre-segmented contours, while the second is based on the differential analysis of the raw grey-scale image. However, in both cases it is the rate of

change of contour angle that is used to characterise corner features.

1.1. Related literature

In the case of boundary-based corner detection from pre-segmented contours there are three processing steps. Firstly, the image is pre-segmented. Secondly, boundaries of the object in the image are extracted and chain-coded. Finally, algorithms are developed for identifying corners in the chain-codes. A common technique is to search for corners at the intersection points or junction points between straight line segments (Xie et al., 1993). In fact, the use of chain-codes to provide a digital characterisation of corners abound in the literature (Freeman and Davis, 1977; Beus and Tiu, 1987; Koplowitz and Plante, 1995; Rosenfeld and

* Corresponding author. Tel.: +44 1904 43 2767; fax: +44 1904 43 2767; e-mail: erh@minster.cs.york.ac.uk

Johmston, 1973; Rosenfeld and Weszka, 1975). A good review is provided by Liu and Srinath (1990a). However, it must be stressed that the weak link in the contour-based method of corner detection is the prior availability of a reliable image segmentation.

Grey-scale corner detection algorithms can be divided into two groups. Template-based corner detectors (Rangarajan et al., 1989; Mehrotra et al., 1990) exploit the similarity between a given template of specific orientation for each image sub-window. Because multiple orientation templates are used, the technique is computationally expensive. Gradient-based corner detectors (Kitchen and Rosenfeld, 1982; Singh, 1990; Deriche and Giraudon, 1993; Noble, 1988; Wang and Brady, 1995; Harris and Stephens, 1988), on the other hand, rely on measuring the curvature of an edge that passes through a given image neighbourhood. The strength of the corner response depends on both the edge-strength and the rate of change of edge-direction. Gradient-based corner detection techniques are more likely to respond to noise than their contour-based counterparts, and often perform quite poorly.

Focusing in more detail on gradient-based corner detection, Wang and Brady (1995) describe a curvature-based corner detector which measures the strength of corners in terms of the so-called “total curvature”. The total curvature is proportional to the second derivative of the image intensity along the edge-tangent direction, and is inversely proportional to the edge-strength. The method offers the attractive feature of exploiting an explicit measure of curvature as well as providing a degree of false corner suppression. However, one of its weaknesses is that it cannot control false responses when significant image noise is present. The popular Plessey operator (Harris and Stephens, 1988) is another curvature-based corner detector which is based only on first derivatives of the image intensity. Although it is quite robust with respect to noise, it can only perform well on ‘L’-junctions or right-angle corners. The localisation of other junction types is poor. The SUSAN corner detector (Smith and Brady, 1997), on the other hand, performs a local template comparison. A degree of robustness is achieved by computing

template correlation statistics for a circular mask centred on the corner candidates. In particular, it examines the number of pixels within the mask which have a similar brightness to the template. The number of agreement counts is taken as a vote for the corner hypothesis. High voting pixels are taken to have a significant corner strength. The method also performs a non-maximum suppression test to reject false-positives.

1.2. Paper outline

Our aim in this paper is to present a new corner detection method which exploits both template-based and gradient-based concepts. We appeal to a gradient-based image representation which has already been shown to provide a convenient topographic representation for edge and symmetry features (Cross and Hancock, 1997, 1998). The representation commences from the Canny edge-map, i.e. the directional gradient of the image intensity. By rotating the edge-gradient vectors we compute a field of edge-tangents. The tangent vectors are smoothed by performing an averaging process in which the vectors are weighted according to the inverse of their distance from the image point in question. This weighting function is suggested by magneto-statics. As a result, we refer to the resulting gradient field as the vector-potential. According to this representation, local maxima of the magnitude vector-field which exhibit directional consistency (i.e. ridges) correspond to edges. Directionally consistent local minima (i.e. ravines) are symmetry lines. In this paper, we extend this topographic picture to include corners. These are saddle-structures where directionally consistent ridges and ravines intersect.

It is important to contrast our method of analysis with Haralick’s topographic primal sketch (Haralick et al., 1983). Whereas Haralick focuses on the topographic structure of grey-scale features (i.e. a scalar image representation), we analyse the topography of a vector-field representation. The main advantage is that we are able to exploit directional consistency in the localisation of topographic structure and, hence, improve the robustness of feature detection. With the topographic representation in hand, the main practical

problem that confronts us is the localisation of the saddle-points. This is more difficult than localising ridges and ravines since we are concerned with identifying point features rather than contour features. In the case of ridges and ravines, we can exploit constraints on compatible continuity or directionality. In the case of saddle-points the constraints are more subtle, since we are seeking locations which are consistent with being the junctions between saddle-ridges and saddle-valleys.

Based on this observation, we develop topographic tests for the consistent saddle-structure in the vector-potential. This is effectively a template-based method. We search for consistent valley structure in the direction of the vector-potential and consistent ridge structure in the orthogonal direction. In other words, the directional template characterises local saddle-structure as the intersection of ridge (i.e. edge) and ravine (i.e. symmetry) structures. Moreover, computation is simplified since the template is fixed to be in the direction of the vector-potential. In this way we avoid explicit computation of directional second derivatives of the image intensity. As we will demonstrate experimentally, this offers advantages in terms of improved noise sensitivity.

The outline of this paper is as follows. In Section 2, we review the vector-potential representation. Section 3 introduces the topographic representation of features in the vector-potential. In Section 4, we confront some of the practical difficulties associated with saddle-localisation. Real-world experimental examples are presented in Section 5. Questions of algorithm sensitivity and comparison are addressed in Section 6. Finally, Section 7 provides some conclusions and identifies avenues for future investigation.

2. Image representation using vector-potential

In this section, we review the feature-representation recently reported by Cross and Hancock (1997). The starting point is to compute the Canny edge-map (Canny, 1986). Accordingly, we commence by convolving the raw image I with a

Gaussian kernel of width σ . The kernel takes the following form:

$$G_\sigma(x, y) = \frac{1}{2\pi\sigma^2} \exp\left[-\frac{x^2 + y^2}{2\sigma^2}\right]. \quad (1)$$

With the filtered image in hand, the Canny edge-map is recovered by computing the gradient

$$\underline{E} = \nabla G_\sigma * I. \quad (2)$$

In order to compute a vector-field representation of the edge-map, we will need to introduce an auxiliary z dimension to the original x - y co-ordinate system of the plane image. In this augmented co-ordinate system, the components of the edge-map are confined to the image plane. In other words, the edge-vector at the point $(x, y, 0)$ on the input image plane is given by

$$\underline{E}(x, y, 0) = \begin{pmatrix} \frac{\partial G_\sigma * I(x, y)}{\partial x} \\ \frac{\partial G_\sigma * I(x, y)}{\partial y} \\ 0 \end{pmatrix}. \quad (3)$$

For an ideal step-edge, the resulting image gradient will be directed along the boundary normal. A more convenient representation is the edge-tangent vectors which flow along the object boundary. Accordingly, we re-direct the edge-vectors so that they are tangential to the original planar shape by computing the cross-product with the normal to the image plane $\hat{z} = (0, 0, 1)^T$. The tangent vector at the point $(x, y, 0)$ on the input image plane is defined to be

$$\underline{j}(x, y, 0) = \hat{z} \wedge \nabla G_\sigma * I(x, y). \quad (4)$$

To be more explicit, in terms of its components the tangent vector is given by

$$\underline{j}(x, y, 0) = \begin{pmatrix} -\frac{\partial G_\sigma * I(x, y)}{\partial y} \\ \frac{\partial G_\sigma * I(x, y)}{\partial x} \\ 0 \end{pmatrix}. \quad (5)$$

In our previously reported work, the key idea underlying the image representation was to characterise edges and symmetry lines using topographic structures in the edge-tangent field. Edges corresponded to locations where the tangent vectors re-enforce one-another. In other words, the boundaries are identified as local maxima of the tangent field. Symmetry points are those at which

there is cancellation between diametrically opposed tangent vectors. Axes of symmetry are lines of local minimum in the tangent field. At the level of fine detail, intensity ridges or ravines give rise to local symmetry axes.

Unfortunately, since the raw gradient vectors are likely to be noisy we must develop a means of smoothing the tangent field so that we can perform the required topographic analysis. To realise this goal, we appeal to magneto-statics to develop a means of smoothing the tangent field. Accordingly, we compute an analogue of the vector-potential by regarding the edge-tangents as a field of elementary currents. The vector potential is found by integrating over volume and weighting the contributing currents according to inverse distance. In other words, the vector-potential at the point $\underline{r} = (x, y, z)^T$ in the augmented space in which the original image plane is embedded is

$$\underline{A}(x, y, z) = \mu \int_{V'} \frac{\underline{j}(x', y', z')}{|\underline{r} - \underline{r}'|} dV', \quad (6)$$

where $\underline{r}' = (x', y', z')^T$ and μ is the permeability constant which we set equal to unity. Since the contributing edge-tangent vectors (or currents) are distributed only on the image plane, the volume integral reduces to an area integral over the image plane. As a result, the components of the vector-potential are as follows:

$$\underline{A}(x, y, z) = \begin{pmatrix} - \int \int \frac{\partial G_{\sigma^* I}(x', y')}{\partial y'} \frac{1}{\sqrt{(x-x')^2 + (y-y')^2 + z^2}} dx' dy' \\ \int \int \frac{\partial G_{\sigma^* I}(x', y')}{\partial x'} \frac{1}{\sqrt{(x-x')^2 + (y-y')^2 + z^2}} dx' dy' \\ 0 \end{pmatrix}. \quad (7)$$

The structure of the vector-potential deserves further comment. In the first instance, the components are confined to the x - y plane for all values of the auxiliary co-ordinate z . However, as we move away from the image plane the role of this auxiliary dimension is to average the generating currents over an increasingly large area of the original image plane. In other words, the role of the auxiliary z -dimension is to allow us to perform volume integration of the edge-tangent vectors. By

sampling the vector-potential for various x - y planes at increasing sampling height z above the image plane, we induce a scale-space representation. We exploit this property to produce a fine-to-coarse image representation as we sample the vector-potential at increasing sampling heights above the physical image plane.

In order to develop the appropriate differential operators for feature characterisation from the vector-potential we have taken the magneto-static analogy one step further and have appealed to the geometry of the associated magnetic field. According to magneto-statics, the magnetic field is the curl of the vector-potential. It is important to stress that because it is less computationally tractable than the vector-potential, the magnetic field is never used directly in our image representation. The role of the magnetic field is to provide an auxiliary representation. The geometry of the magnetic field allows us to understand the differential structure of the vector-potential. According to our representation of the image structure, symmetry lines follow the local minima of the vector-potential. In other words, they connect image points where there is strong cancellation edge tangent vectors associated with symmetrically placed object boundaries. By contrast, edge-contours follow the local maxima of the vector-potential. According to our representation, the edge-lines connect points where there is strong directional re-enforcement between edge-tangent vectors. Symmetry lines can be interpreted as locations where the magnetic field is perpendicular to the sampling image plane. Edges are locations where field lines are tangential to the relevant sampling plane. When viewed from the perspective of the differential structure of the vector-potential, symmetry lines are locations where the component of the curl in the image plane vanishes, i.e. $\hat{z} \wedge \nabla \wedge \underline{A}(x, y, z) = 0$, edges are locations where the transverse component of the divergence vanishes, i.e. $\nabla \cdot (\hat{z} \wedge \underline{A}(x, y, z)) = 0$.

Corners, or points of locally maximum boundary curvature, can be viewed as edge-locations where there is a local symmetry axis associated with a rapid change in boundary direction. When viewed from the perspective of our image representation, corners therefore correspond to

locations where both the edge and symmetry conditions are simultaneously satisfied. From a topographic viewpoint, corners are located where boundary lines and symmetry lines meet. In other words, we are interested in locating points where there is a local maximum of the magnitude of the vector-potential in one direction and a local minimum in the orthogonal direction. As a result, corner detection can be treated as saddle-point detection.

3. Topographic representation

In Section 2, we established that corners are saddle-points in the magnitude of the vector-potential. We therefore focus on the analysis of the scalar quantity

$$g(x, y)_z = |\underline{A}(x, y, z)|. \tag{8}$$

The topographic structure of the vector-potential can be characterised using the Hessian matrix

$$\mathcal{H}_g = \begin{pmatrix} g_{xx} & g_{xy} \\ g_{xy} & g_{yy} \end{pmatrix}, \tag{9}$$

where the second derivatives are given by

$$g_{xx} = \frac{\partial^2 g(x, y)_z}{\partial^2 x},$$

$$g_{xy} = \frac{\partial^2 g(x, y)_z}{\partial x \partial^2 y},$$

$$g_{yy} = \frac{\partial^2 g(x, y)_z}{\partial^2 y}.$$

The eigen-structure of the Hessian matrix can be used to gauge the curvature of the surface. The two eigen-values of \mathcal{H} are the maximum and minimum curvatures. The orthogonal eigen-vectors of \mathcal{H} are known as the principal curvature directions. The mean-curvature (K) of the surface is found by averaging the maximum and minimum curvatures. Finally, the Gaussian curvature (H) is equal to the product of the two eigen-values. As a result,

$$H = g_{xx}g_{yy} - g_{xy}^2 \tag{10}$$

and

$$K = \frac{g_{xx} + g_{yy}}{2}. \tag{11}$$

Table 1
Curvature classes

Class	Symbol	K	H	Region-type
Dome	D	–	+	Elliptic
Ridge	R	–	0	Parabolic
Saddle-ridge	SR	–	–	Hyperbolic
Plane	P	0	0	Hyperbolic
Saddle-point	S	0	–	Hyperbolic
Cup	C	+	+	Elliptic
Valley	V	+	0	Parabolic
Saddle-valley	SV	+	–	Hyperbolic

The signs and zeros of the mean and Gaussian curvatures can be used to categorise the local surface geometry into a number of distinct topographic classes. These classes are summarised in Table 1. In this paper, we are interested in the saddle-structures which are labelled as hyperbolic features in the table. These features are characterised by the condition $H < 0$. In particular, we are interested in points that are consistent with being the intersections of edge and symmetry lines, i.e. in the intersections of saddle-ridges and saddle-valleys. The joint condition for the intersections is

$$K \neq 0 \wedge H < 0. \tag{12}$$

By searching for the intersection of consistent saddle-ridges, we overcome some of the problems of localising saddle-points, for which $K = 0$ and $H < 0$. This can prove difficult since there are no constraints from the directionality of the desired feature. In particular, we mitigate this difficulty and realise the corner localisation process using templates to search for the junctions between symmetry lines and edge-lines.

4. Implementation


In this section, we describe two aspects of the implementation of our corner detector. The first of these is the means by which we compute the vector-potential and the resulting computational complexity of the method. The computation is realised using fast Fourier transforms. The second

implementational detail concerns the practical means by which we locate saddle-structures.


4.1. Computing the vector-potential

Key to our implementation is the fact that the volume integrals appearing in the definition of the


vector-potential (Eq. (7)) can be replaced by spatial-convolutions with a sampling height-dependent filter. Specifically, we invoke the Fourier duality between convolution in the spatial domain and multiplication in the frequency domain. In this way, the discretised version of the vector-potential can be computed using just three 2D Fourier



(a) Magnitude



(b) Direction




(c) Detected corners

Fig. 1. Topographic representation and corner detection.


transforms and a pair of frequency-domain convolutions.

Our basic goal is to compute the vector-potential at a given sampling height above the image plane. The 2D integrals appearing in the definition of the vector-potential can be discretised to give the following x and y components:


$$A_x(x, y, z) = - \sum_{x'} \sum_{y'} \frac{\partial G_\sigma * I(x', y')}{\partial y'} \times \frac{1}{\sqrt{(x-x')^2 + (y-y')^2 + z^2}}, \quad (13)$$




(a) Magnitude at sampling height $z = 0$



(b) Direction at sampling height $z = 0$



(c) Magnitude at sampling height $z = 1$




(d) Direction at sampling height $z = 1$

Fig. 2. Topographic representation at different sampling heights.


$$A_y(x, y, z) = \sum_{x'} \sum_{y'} \frac{\partial G_\sigma * I(x', y')}{\partial x'} \times \frac{1}{\sqrt{(x-x')^2 + (y-y')^2 + z^2}}. \quad (14)$$

The double summation can be replaced by a convolution with a composite filter. For instance, the x -component of the vector-potential is as follows:


$$A_x(x, y, z) = -(V_x(\sigma, z) * I)(x, y). \quad (15)$$




(a) Magnitude at sampling height $z = 2$



(b) Direction at sampling height $z = 2$



(c) Magnitude at sampling height $z = 3$



(d) Direction at sampling height $z = 3$

Fig. 3. Topographic representation at different sampling heights.

We exploit the commutative properties of convolution to compute the composite filter $V_x(\sigma, z)$. The filter is found by convolving the appropriate directional derivative of the Gaussian with the inverse Euclidean distance operator, i.e.,

$$\begin{aligned}
 &V_x(\sigma, z)(x, y) \\
 &= \sum_{x'} \sum_{y'} \frac{\partial G_\sigma(x', y')}{\partial y'} \\
 &\quad \times \frac{1}{\sqrt{(x-x')^2 + (y-y')^2 + z^2}}. \tag{16}
 \end{aligned}$$

For practical reasons, we would like to realise the computation of the components of the vector-potential using fast Fourier transforms. Our basic strategy is to exploit the Fourier duality between convolution in the spatial-domain and multiplication in the frequency domain. Schematically, we utilise the identity

$$A_x = \mathcal{F}^{-1}[\mathcal{F}[V_x(\sigma, z)] \times \mathcal{F}[I]] \tag{17}$$

to compute each of the components of the vector-potential in turn. In this way the vector-potential can be obtained using three separate Fourier transform operations. The first of these involves computing the Fourier transform of the raw image $\mathcal{F}[I]$. Two separate weighted spatial frequency distributions are then constructed by multiplying the components of the image Fourier transform with the Fourier representation for each of the

composite filters in turn, i.e. we compute $\mathcal{F}[A_x(\sigma, z)]$ and $\mathcal{F}[A_y(\sigma, z)]$. Finally, the two spatial components of the vector-potential are obtained by inverse Fourier transformation of the two weighted frequency distributions.

Since the computation is implemented using fast Fourier transforms, the time complexity is of the order $(N \ln(N))^2$ where N is the linear image dimension. In fact, we have implemented the algorithm on an SGI Indy workstation where it is capable of processing 256×256 pixel images at the rate of 25 frames per second.

4.2. Localising saddles

Based on the results presented in Section 3, we make the following observations concerning the topographic structure of the vector-potential in the proximity of corners:

- There is a local minima of the magnitude of the vector-potential in the direction of the vector-potential.
- There is a local maxima of the magnitude of the vector-potential in the orthogonal direction.
- At the locations of corners, the magnitude of vector-potential along both the contour and its orthogonal direction changes rapidly.
- At the locations of corners, the magnitude of the Gaussian curvature is significant.

Based on the first two observations, we search for saddle-points that are consistent when viewed



(a) Original image

(b) Corners from new method

(c) Corners from SUSAN

Fig. 4. Corner detection results.

from a finite support neighbourhood. In practice we localise consistent topographic structure using a simplified form of template convolution. Our template tests for orthogonal maxima and minima

using directional second derivatives and subsequent non-maximum suppression and non-minimum suppression tests. The saddle-points are corner candidates. Because of image noise and



Fig. 5. Corner detection at different sampling heights.

other imperfections, the points detected by our saddle-template are not always the locations of true corners in the image. To overcome this problem we can appeal to the directional consistency of the derivatives of the vector-potential to refine the corner estimates. The aim is to search for orthogonal ridge and ravine structures. To meet this goal, we used directional second derivative operators to compute a corner “strength” measure. This measure captures the directional variations of the vector-potential along the contour direction and in the orthogonal direction.

To be more formal, suppose that $V(x, y, z) = \nabla_{\parallel}^2 |\vec{A}|$ is the second derivative of the magnitude of the vector-potential in the direction of the vector-potential. The magnitude of this quantity will take on a maximum value when there is a local minimum or valley structure in the magnitude of the vector-potential. Further suppose that $R(x, y, z) = \nabla_{\perp}^2 |\vec{A}|$ is the second derivative of the magnitude of the vector-potential in the direction perpendicular to the local orientation of the vector-field. The magnitude of this quantity exhibits a local maxima when there is a ridge structure or local maxima in the vector-potential. Using these two operators, we search for corners by computing the following corner strength measure which gauges consistent saddle-structure:

$$C = |V(x, y, z)| \times |R(x, y, z)|. \quad (18)$$

The measure is an approximation to the Gaussian curvature. It is large in value when there are orthogonal ridges and valleys in the magnitude of the vector-potential. Corners are selected by

thresholding this aggregate measure of corner strength.

5. Experiments

In this section, we provide some experimental evaluation of the corner detection algorithm. The experimental work is divided into two parts. We commence with some examples on binary imagery to illustrate some of the properties of the representation. Next we furnish real-world examples. To illustrate the properties of our vector-potential representation and corner detection algorithm, we use a simple binary image of “E”. Fig. 1(a) is the magnitude of the vector-potential for the binary image, Fig. 1(b) is the direction of the vector-potential. Fig. 1(c) shows the detected corners. For this simple image, the results are all correct. The magnitude of the vector-potential is displayed as a function of the sampling height z in Fig. 1(a) to emphasise the topographic structure. Here the saddle-structure associated with the corners is clear. The ridge and ravine structure of the edge and symmetry lines is also evident. In Fig. 1(b) we display the vectorial representation of $\underline{A}(x, y, 0)$. The main feature to note from this figure is that the direction of the vector-potential changes rapidly at the corner locations.

As explained earlier, we can endow our image representation with a scale-space dimension by sampling the vector-potential at increasing sampling heights above the image plane. In Figs. 2 and 3, we provide some qualitative examples of this




Fig. 6. Cog-wheel test images.






Fig. 7. Comparison for different sampling heights.

scale-space sampling. In each case the left-hand figure is an elevation map showing the magnitude of the vector-potential while the right-hand figure is the vector-field. The main feature to note from these examples is that as the scale or sampling height is increased, so the saddle-structures become shallower.


We now turn our attention to real-world scenes. To provide some comparison, we have provided some experimentation with the SUSAN corner detector (Smith and Brady, 1997). Fig. 4(a) is the original INRIA office image. In Fig. 4(b) we show the result of applying the algorithm reported in this paper, while Fig. 4(c) shows the result of ap-




(a) Angle = 30 degrees



(b) Angle = 60 degrees



(c) Angle = 90 degrees



(d) Angle = 120 degrees

Fig. 8. Saddle-structures for different opening angles.

plying the SUSAN corner detector. The results obtained with our algorithm are generally cleaner, and there are fewer false positives. There are also some interesting qualitative differences in the detected corners. For instance in our algorithm, the meeting of the line-like horizontal bars and thicker vertical bars of the window are detected as single junctions. In the case of SUSAN, double corners are returned. The result of our algorithm is more perceptually intuitive and may prove more useful for higher level matching problems.

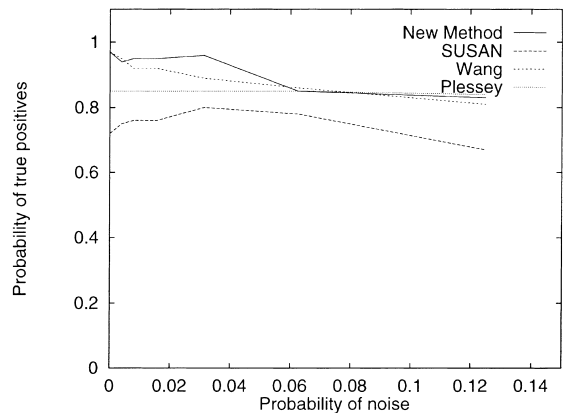
Finally, we provide some examples of the scale-space detection of corners in the INRIA office scene. Fig. 5 shows the results of corner detection at a number of different sampling heights. The left column of the figure shows the magnitude of the vector-potential, while the right-column shows the detected corners superimposed on the original image. As we move from the top row of the figure to the bottom row, the sampling height z of the vector-potential increases. As the sampling height increases, then so only the dominant corners remain. However, the majority of the significant corners persist over the full set of sampling heights.

6. Performance analysis

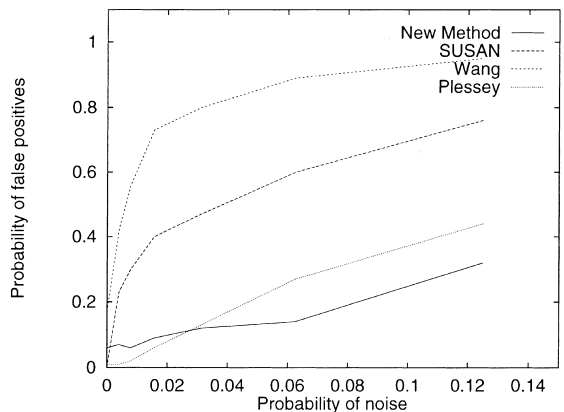
Our final piece of experimental work is aimed at measuring the noise sensitivity of our corner detector and comparing it with some of the alternative corner detectors reviewed in the introductory section of this paper. The specific algorithms used in this comparison are the SUSAN corner detector (Smith and Brady, 1997), Wang and Brady’s corner detector (Wang and Brady, 1995) and the Plessey corner detector (Harris and Stephens, 1988). To realise this comparison, we have generated synthetic images of cog-wheels (see Fig. 6) and have added salt-and-pepper noise with known proportion. By increasing the number of spikes on the circumference of the cog, we can systematically reduce the opening angle of the corners. The synthetic figure provides ground-truth data in which the number of target corners is known. We focus on two aspects of the noise sensitivity of our corner detector. The first of these is the scale-depend-

dance of the corner detection process. The second is the error rate for false positives and false negatives.

We commence our evaluation by measuring the accuracy of the corner detection process as a function of sampling height (i.e. spatial scale) and corner opening angle. Fig. 7 shows the fraction of correctly detected corners as a function of sampling height z . The different curves are for different opening angles. The main conclusion that can be drawn from this plot is that our corner detector degrades with increasing sampling height. There is also an indication that we encounter difficulties with small opening angle corners.



(a) True Positives




(b) False Positives


Fig. 9. Comparison of noise sensitivity.

In order to understand in a qualitative way the angle systematics involved in corner detection, we have generated a series of synthetic examples. The resulting plots of the vector-potential magnitude


are shown in Fig. 8. As the angle increases, so the depth of the saddle decreases. At small opening angles, the width of the saddle becomes very narrow and hence difficult to localise.




(a) True Positives, Opening angle = 30 degrees




(b) False Positives, Opening angle = 30 degrees




(c) True Positives, Opening angle = 90 degrees



(d) False Positives, Opening angle = 90 degrees



(e) True Positives, Opening angle = 120 degrees



(f) False Positives, Opening angle = 120 degrees

Fig. 10. Noise sensitivity for various opening angles.

The second aspect of our sensitivity study is to provide some comparison with the alternative corner detectors under conditions of controlled noise. The plots in Fig. 9 show the probability of both true positives, i.e. the fraction of genuine corners that are correctly detected, and the probability of false positives, i.e. the fraction of detected corners which are false alarms. The probabilities of true positives and false positives are plotted as a function of the probability or fraction of added noise. In each case, the solid curve is the result of the proposed algorithm, the long dashed curve is for the SUSAN corner detector, the dashed curve is for Wang and Brady's corner detector and the dotted line is for the Plessey corner detector. From the plot of Fig. 9(a), we can draw the conclusion that the proposed algorithm consistently outperforms the alternatives in the sense that it has a higher probability of true positives for small portions of added noise. When more noise is added to the image, the performance of our method is almost identical with that of the Wang and Brady, and the Plessey corner detectors. Our method always outperforms SUSAN in terms of its probability of true positives.

The results shown in Fig. 9(b) provide more conclusive support for our proposed corner detector. Here we see that the probability of false positives for our proposed method is consistently significantly lower than that for the alternative corner detectors. The only exception occurs at low noise levels where the Plessey corner detector offers slightly better performance.

Next, we aim to compare the performance of the corner detectors for different opening angles and different sampling heights. We choose a small opening angle of 30° , a medium opening angle of 90° and a large opening angle of 120° . From Fig. 10, we draw the conclusion that, when measured in terms of the fraction of true positives, our method performs a little worse than that of the other corner detectors for small opening angles, a little better than the others at medium opening angles, and significantly better than the alternatives at large opening angles. When gauged in terms of the fraction of false positives, our method is better than the alternatives at both small opening angles and large

opening angles. However, it performs a little worse than the Plessey corner detector at medium opening angle.

7. Conclusions

In this paper, we have presented a corner detection method that is based on the topographic analysis of a vector-potential image representation. According to our representation, corners are saddle-points where saddle-ridges and saddle-valleys intersect. Experimental results reveal the method offers performance advantages over the SUSAN corner detector, Wang and Brady's corner detector and the Plessey corner detector. The most striking of these is to offer better control over false positives.

There are a number of ways in which the ideas presented in this paper could be developed. We clearly have a means of computing a curvature scale-space. In this respect, our work is similar to that of Asada and Brady (1986) and Mokhtarian and Mackworth (1992). Our next step is to emulate this work by investigating how the new corner representation can be used for shape matching and recognition.

References

- Asada, H., Brady, M., 1986. The curvature primal sketch. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 8 (1), 2–14.
- Beus, H.L., Tiu, S.S.H., 1987. An improved corner detection algorithm based on chain-code plane-curves. *Pattern Recognition* 20 (20), 291–296.
- Canny, J.F., 1986. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 8 (6), 679–698.
- Costabile, M.F., Guerra, C., Pieroni, G., 1985. Matching shapes: A case study in time varying images. *Computer Vision, Graphics and Image Processing* 29, 296–310.
- Cross, A.D.J., Hancock, E.R., 1997. Scale-space vector field for feature analysis. In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, San Juan, Puerto Rico. IEEE Computer Society, Silver Spring, MD, pp. 738–743.
- Cross, A.D.J., Hancock, E.R., 1998. Scale-space vector field for symmetry detection. *Image and Vision Computing* 11, to appear.

- Deriche, R., Giraudon, G., 1993. A computational approach for corner and vertex detection. *International Journal of Computer Vision* 10, 101–124.
- Dreschler, L., Nigel, H., 1992. Volumetric model and 3D trajectory of a moving car derived from monocular TV-frame sequence of a street scene. In: *Proceedings of IJCAI* pp. 692–697.
- Freeman, H., Davis, L.S., 1977. A corner finding algorithm for chain code curves. *IEEE Transactions on Computers* 26, 297–303.
- Han, M.H., Jang, D., 1990. The use of maximum curvature points for the recognition of partially occluded objects. *Pattern Recognition* 23, 21–33.
- Haralick, R., Watson, L., Laffey, T., 1983. The topographic primal sketch. *IJRR* 2, 50–72.
- Harris, C.G., Stephens, M., 1988. A combined corner and edge detector. In: *4th Alvey Vision Conference*, pp. 147–151.
- Kitchen, L., Rosenfeld, A., 1982. Gray-level corner detection. *Pattern Recognition Letters* 1, 95–102.
- Koplowitz, J., Plante, S., 1995. Corner detection for chain-coded curves. *Pattern Recognition* 28 (6), 843–852.
- Liu, H.C., Srinath, M.D., 1990a. Corner detection from chain-code. *Pattern Recognition* 23 (1–2), 51–68.
- Liu, H.C., Srinath, M.D., 1990b. Partial classification using contour matching in distance transformation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12, 1072–1079.
- Mehrotra, R., Nichani, S., Ranganathan, N., 1990. Corner detection. *Pattern Recognition* 23 (11), 1223–1233.
- Mokhtarian, F., Mackworth, A.K., 1992. A theory of multi-scale curvature-based shape representation for planar curves. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14 (8), 789–805.
- Noble, J.A., 1988. Finding corners. *IVC* 6 (2), 121–128.
- Rangarajan, K., Shah, M., Brackle, D.V., 1989. Optimal corner detection. *Computer Vision Graphics and Image Processing* 48, 230–245.
- Rosenfeld, A., Johnstone, E., 1973. Angle detection on digital curves. *IEEE Transactions on Computers* 22, 875–878.
- Rosenfeld, A., Weszka, J.S., 1975. An improved method of angle detection on digital curves. *IEEE Transactions on Computers* 24, 940–941.
- Singh, A., 1990. Gray level corner detection – A generalization and a robust real-time implementation. *Computer Vision Graphics and Image Processing* 51 (1), 54–69.
- Smith, S., Brady, J., 1997. SUSAN – A new approach to low level image processing. *International Journal of Computer Vision* 23 (1), 45–78.
- Wang, H., Brady, M., 1995. Real-time corner detection algorithm for motion estimation. *IVC* 13 (9), 695–703.
- Xie, X., Sudhakar, R., Zhuang, H., 1993. Corner detection by a cost minimization approach. *Pattern Recognition* 26 (8), 1235–1243.