Application of the Fisher-Rao Metric to Ellipse Detection

STEPHEN J. MAYBANK

School of Computer Science and Information Systems, Birkbeck College, Malet Street, London, WC1E 7HX, UK.
sjmaybank@dcs.bbk.ac.uk

Abstract. The parameter space for the ellipses in a two dimensional image is a five dimensional manifold, where each point of the manifold corresponds to an ellipse in the image. The parameter space becomes a Riemannian manifold under a Fisher-Rao metric, which is derived from a Gaussian model for the blurring of ellipses in the image. Two points in the parameter space are close together under the Fisher-Rao metric if the corresponding ellipses are close together in the image. The Fisher-Rao metric is accurately approximated by a simpler metric under the assumption that the blurring is small compared with the sizes of the ellipses under consideration. It is shown that the parameter space for the ellipses in the image has a finite volume under the approximation to the Fisher-Rao metric. As a consequence the parameter space can be replaced, for the purpose of ellipse detection, by a finite set of points sampled from it. An efficient algorithm for sampling the parameter space is described. The algorithm uses the fact that the approximating metric is flat, and therefore locally Euclidean, on each three dimensional family of ellipses with a fixed orientation and a fixed eccentricity. Once the sample points have been obtained, ellipses are detected in a given image by checking each sample point in turn to see if the corresponding ellipse is supported by the nearby image pixel values. The resulting algorithm for ellipse detection is implemented. A multiresolution version of the algorithm is also implemented. The experimental results suggest that ellipses can be reliably detected in a given low resolution image and that the number of false detections can be reduced using the multiresolution algorithm.

Keywords: ellipse detection, Fisher-Rao metric, flat metric, geodesic, Hough transform, Kullback-Leibler distance, lattice, multiresolution, Riemannian manifold, volume of a Riemannian manifold, Voronoi's principal lattice.

1   Introduction

One of the main tasks in computer vision is the detection of structures in images. Examples of image structures include simple curves such as lines, circles and ellipses (Fitzgibbon et al., 1999; Kanatani, 1994, 1996; Leedan and Meer, 1998; Maybank, 2004; Rosin, 1996; Taubin, 1991; Zhang, 1997) as well as more abstract structures such as projective transformations of the line, collineations, epipolar transforms and fundamental matrices (Hartley and Zisserman, 2003; Leedan and Meer, 1998; Maybank, 2003, 2005; Torr and Fitzgibbon 2004). In many cases the structures of interest form a parameterised family. For example, the family of all lines in a two dimensional image can be parameterised by the points of

a two dimensional manifold and the family of all ellipses in a two dimensional image can be parameterised by the points of a five dimensional manifold.

A general theory for structure detection is developed in Maybank (2004, 2006). The image, or more precisely the image domain, is regarded as a Riemannian manifold, D, and the structures are subsets of D. In most cases D is an open subset of a Euclidean space and the metric on D is the Euclidean metric. The probability of obtaining a measurement x in D, given that the structure θ is present, is described by a conditional probability density function $p(x|\theta)$. The conditional density, $p(x|\theta)$, is obtained by blurring an initial density concentrated on the structure θ in D. The resulting density depends on the metric defined on D as well as on the structure θ.

Let T be a parameter space for the family of structures. The family $\theta \mapsto p(x|\theta)$, $\theta \in T$ of conditional densities defines a Riemannian metric on T known as the Fisher-Rao metric $J(\theta)$ (Amari, 1985; Cover and Thomas, 1991; Maybank, 2003, 2004). In this context, $J(\theta)$ is the only correct metric for measuring the distances between nearby points of T (Amari, 1985). This is because $J(\theta)$ determines the leading order approximation to the expected log likelihood ratio $\ln(p(x|\theta)/p(x|\theta'))$ where θ' is a point of T near to θ,

$$\int_D \ln\left(\frac{p(x|\theta)}{p(x|\theta')}\right) p(x|\theta)\, dx = \frac{1}{2}(\theta - \theta')^\top J(\theta)(\theta - \theta') + O_3, \tag{1}$$

where O_3 consists of terms third order or higher in $\theta - \theta'$. If two structures θ, θ' are close together under $J(\theta)$, then it is difficult to distinguish between them given a measurement x. If $p(x|\theta)$ is large, indicating that x is near to the structure θ, then with a high probability, $p(x|\theta')$ is also large, indicating that x is near to the structure θ'.

In most cases there is no closed form expression for the Fisher-Rao metric, however there exists an asymptotic approximation which is accurate if the noise level is low. In favourable cases such as lines (Maybank 2004) and projective transformations of the line (Maybank 2005) the approximating metric can be expressed in terms of standard functions. It is shown below in Section 3.3 that ellipses are included amongst the favourable cases.

The Fisher-Rao metric, or in practice the approximating metric, is the basis of a structure detection algorithm in which T is sampled at a finite number of points $\theta(i)$, $1 \leq i \leq n_s$, and each $\theta(i)$ is checked to see if the corresponding structure is present in D. The minimum number n_s of sample points depends on the volume of T under the Fisher-Rao metric. The advantage of this sample based approach to structure detection is that it is extremely robust in the presence of outliers, i.e. those measurements which do not provide any information about the particular structure $\theta(i)$ which is being checked. All the measurements outside a small neighbourhood of the image structure $\theta(i)$ can be discarded. Once the structure is detected, the estimate $\theta(i)$ of the parameter vector associated with the structure can be improved using standard methods.

1.1 Contributions

The main contributions of this work are:

(i) a computationally tractable approximation $K(\theta)$ to the Fisher-Rao metric on the parameter space for ellipses in two dimensional images;

2

(ii) a proof that the three dimensional parameter space for the set of ellipses with a given orientation and a given eccentricity is flat under $K(\theta)$;

(iii) a new algorithm for choosing a finite set of sample points with which to approximate the parameter space T for the ellipses;

(iv) an ellipse detection algorithm in which each sample point in T is checked to see if the corresponding ellipse is supported by the values of the image pixels near to it.

Contribution *(i)* provides answers to two long standing questions about the Hough transform for ellipses: how large should the accumulators be? and how many accumulators are required? For further information see the remarks at the end of Section 4.2.

1.2   Summary by section

Related work on ellipse detection is described in Section 2. Families of plane curves, their associated Fisher-Rao metrics, and the family of ellipses in particular are discussed in Section 3. The algorithm for sampling the parameter space for ellipses is described in Section 4. The criteria for deciding if the presence of a given ellipse is supported by the image grey levels are described in Section 5 and experimental results are presented. Some suggestions for future work are made in Section 6.

2   Related Work

2.1   Detection of ellipses and other structures

The term 'ellipse detection' is used in the current work for cases in which an image is given and it is necessary to decide if the image contains an ellipse. The image may contain several ellipses or alternatively it may contain no ellipses at all. Even when an ellipse is present, large parts of the image, possibly including parts of the image near to the ellipse, may contain no information supporting the presence of the ellipse. Once an ellipse is detected, those parts of the image remote from it can be discarded and the ellipse located accurately using information from those parts of the image near to it. Algorithms for locating ellipses accurately are described by Fitzgibbon et al.(1999), Kanatani (1994, 1996), Leedan and Meer (1998), Rosin (1996), Taubin (1991), Zhang (1997).

A number of Hough transform approaches to ellipse detection are described in Kanatani and Ohta (2003) and Leavers (1992). Applications of RANSAC to ellipse detection are described by Aguado et al. (1995) and the randomised Hough transform, which is a combination of the Hough transform and RANSAC, is described by Xu and Oja (1993).

Yao et al. (2004) and Lutton and Martinez (1994) use genetic algorithms to detect ellipses in images. In Yao et al. (2004) each so called chromosome records the coordinates of five points in the image. The relevant chromosomes are those in which the five points lie on a unique ellipse. In Lutton and Martinez (1994) each chromosome records five parameters which, taken together, uniquely specify an ellipse. The fitness of a chromosome is calculated by examining the image feature points near to the ellipse specified by the chromosome.

Ji and Haralick (1999) detect ellipses by assembling arc segments in an image. The arcs are obtained by chaining the outputs of an edge detector. A statistical criterion is used to compare the parameters of neighbouring arcs in order to to decide if the arcs belong to a single ellipse. Xie and Ji (2002) fit ellipses to sets of points by checking each pair of points in turn to see if the two points are at opposite ends of the major axis of an ellipse whose presence in the image is supported by a sufficiently large subset of the remaining points.

The results obtained in Sections 3 and 4 below show that all ellipse detection methods which rely on the bottom up accumulation of image information or the detection of particular points on ellipses can be subsumed in a straightforward search through a finite set of ellipses obtained by sampling the parameter space T for ellipses.

Arias-Castro et al. (2005) describe a general theory for detecting two dimensional structures in digital images. The pixels outside the structure take white noise values. In their theory the pixels within the structure take values obtained by adding a constant offset to the white noise. Optimal detection methods are described for a range of structures including line segments, rectangles, disks and objects bounded by an ellipse. In their analysis of structure detection, Arias-Castro et al. (2005) define a distance between any two given structures. The distance depends on the degree of overlap of the structures in the image. In contrast with the Fisher-Rao metric, the distance measure used by Arias-Castro et al. is not closely related to the likelihood ratio for comparing different hypotheses about the location of a structure.

Desolneux at al. (2000, 2003) describe a statistical theory of line detection in digital images. A line is detected if there are sufficiently many pixels along the line at which the direction of the line is normal to the grey level gradient, to within prescribed error bounds. This work can be extended from lines to ellipses and other curves. If this extension were to be carried out, then it would lead to new criteria for deciding if the presence of a given ellipse, sampled from T, is supported by the pixel values.

2.2 Applications

The applications of ellipse detection include object location, object recognition, tracking and camera calibration. Maio and Maltoni (1998) use a generalised Hough transform to locate image ellipses which approximate to the boundaries of faces. Wang et al. (2003) fit an ellipse to the image of the outer boundary of the iris in the human eye. Then, under the assumption that the outer boundary of the iris in 3D is a circle, they infer the 3D orientation of the eye. Wang et al. use the fact that if the camera calibration is known there are up to scale only two object circles that can project to the ellipse fitted to the boundary of the image of the iris.

Agrawal and Davis (2003) calibrate a camera by fitting ellipses to the boundaries of three separate images of the same sphere. Scaggiante et al. (1999) use an ellipse as a model for the image of the circular opening of a pipe. The elliptical image is tracked as the camera moves relative to the pipe. An ellipse model of the image of the human head is used by a particle filter head tracker in Kwolek (2004). Lazebnik et al. (2004) use groups of features within elliptical boundaries to describe the local appearance of objects. When matching features they use a distance measure between ellipses, but their distance measure is different from the approximation to the Fisher-Rao metric obtained below in

Section 3.3.

3   Plane Curves Under Affine Transformations

Let $A(2)$ be the group of affine transformations of the plane (Hartley and Zisserman 2003) and let C be a fixed curve in the plane. A parameterised family of curves in the plane is obtained by applying the elements of $A(2)$ to C (Younes 1998). The notation associated with this parameterised family is introduced in Section 3.1, the approximation $K(\theta)$ to the Fisher-Rao metric for this parameterised family is described in Section 3.2. The metric $K(\theta)$ for ellipses in the plane is obtained as a particular application of this theory in Sections 3.3 and 3.4.

3.1   Curve parameterisation

It is assumed that C is a closed curve given in parameterised form by the C^2 functions z_1, z_2 defined on the unit circle,

$$\phi \mapsto z(\phi) \equiv (z_1(\phi), z_2(\phi))^\top, \qquad 0 \leq \phi < 2\pi. \qquad (2)$$

Each point of $A(2)$ acts on C to produce a closed curve in the plane. This yields a family of closed curves parameterised by the points of $A(2)$. In particular applications, including the application to ellipses in Section 3.3, the parameter space may be a proper subset of $A(2)$, rather than the whole of $A(2)$. Let $T \subseteq A(2)$ be the chosen parameter space. It is assumed that T is a manifold although T may not be a submanifold of $A(2)$. Let θ be a parameter vector for T and let the affine transformation defined by θ be

$$z \mapsto L(\theta)z + m(\theta), \qquad (3)$$

where $L(\theta)$ is an invertible 2×2 matrix and $m(\theta)$ is a vector in $\mathbb{R}^2$. The affine transformation of C is the curve $C(\theta)$ parameterised by $\phi \mapsto q(\theta, \phi)$, where

$$q(\theta, \phi) = L(\theta) \begin{pmatrix} z_1(\phi) \\ z_2(\phi) \end{pmatrix} + m(\theta), \qquad 0 \leq \phi < 2\pi. \qquad (4)$$

Let $\|.\|$ be the Euclidean norm in the plane, and let $\dot{z}(\phi)$, $a(\theta, \phi)$ be defined by

$$\begin{aligned} \dot{z}(\phi) &= (dz_1/d\phi, dz_2/d\phi)^\top, & 0 \leq \phi < 2\pi, \\ a(\theta, \phi) &= \|L(\theta)\dot{z}(\phi)\|, & 0 \leq \phi < 2\pi. \end{aligned}$$

The unit tangent vector to $C(\theta)$ at $q(\theta, \phi)$ is $a(\theta, \phi)^{-1} L(\theta)\dot{z}(\phi)$ and the arc length measure ds on $C(\theta)$ is $ds = a(\theta, \phi)\, d\phi$. The unit normal $n(\theta, \phi)$ to $C(\theta)$ at $q(\theta, \phi)$ is

$$n(\theta, \phi) = a(\theta, \phi)^{-1} \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} L(\theta)\dot{z}(\phi), \qquad 0 \leq \phi < 2\pi.$$

The unit tangent vector and the unit normal to $C(\theta)$ at $q(\theta, \phi)$ are illustrated in Fig. 1.
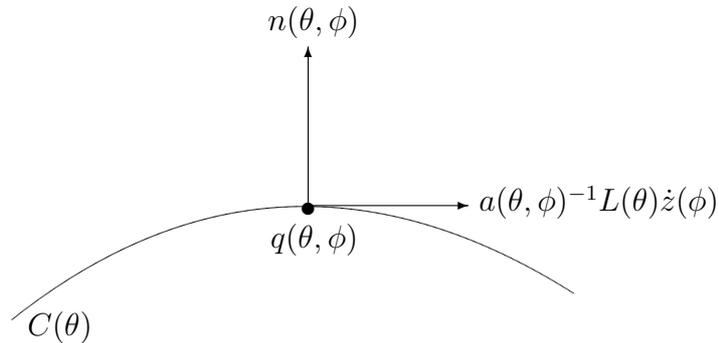
Figure 1. Unit normal and unit tangent vector to the curve $C(\theta)$ at $q(\theta, \phi)$.

3.2 Approximation to the Fisher-Rao metric

The notation ∂ is used for partial derivatives. For example, if $x \mapsto f(x)$ is a real valued function of the vector $x = (x_1, x_2)$, then $\partial_x f(x) = (\partial_{x_1} f(x), \partial_{x_2} f(x))$. It is assumed that the image D is part of the Euclidean plane $\mathbb{R}^2$ and that D carries the Euclidean metric inherited from $\mathbb{R}^2$. An expression is obtained for the conditional probability density function $p(x|\theta)$ referred to in Section 1.

It is assumed that each measurement x arises from an unknown error free measurement $\tilde{x}$ on $C(\theta)$. The probability density function $p(x|\tilde{x})$ is chosen to be Gaussian,

$$p(x|\tilde{x}) = \frac{1}{4\pi t} \exp\left(-\frac{1}{4t}\|x - \tilde{x}\|^2\right), \qquad t > 0. \tag{5}$$

The standard deviation σ of the Gaussian density $p(x|\tilde{x})$ is $\sigma = (2t)^{1/2}$.

The distribution of $\tilde{x}$ on $C(\theta)$ is described by a probability measure dh on $C(\theta)$. The conditional density, $p(x|\theta)$, is obtained by integrating $p(x|\tilde{x})$ over $C(\theta)$, using the measure $dh(\tilde{x})$,

$$p(x|\theta) = \int_{\tilde{x} \in C(\theta)} p(x|\tilde{x})\, dh(\tilde{x}).$$

Let $l(\theta)$ be the total arc length of $C(\theta)$, calculated using the metric induced on $C(\theta)$ by the Euclidean metric on D. The density, dh, for the noise free measurement $\tilde{x}$ on $C(\theta)$ is chosen proportional to the arc length s,

$$dh(s) = l(\theta)^{-1} ds = l(\theta)^{-1} a(\theta, \phi)\, d\phi. \tag{6}$$

The Fisher-Rao metric, $J(\theta)$, (Amari 1985) is defined by

$$J(\theta) = -\int_D (\partial_{\theta\theta}^2 \ln p(x|\theta)) p(x|\theta)\, dx.$$

It is shown in Maybank (2006) that the asymptotic approximation $K(\theta)$ to $J(\theta)$ is

$$K(\theta) = -\int_{C(\theta)} \left(\partial_{\theta'\theta'}^2 \ln p(y|\theta')\right)_{y=\tilde{x}(s), \theta'=\theta} dh(s). \tag{7}$$

6

See also Section 14.4 of Kanatani (1996). Let $w(x, \theta)$ be the minimum distance from the point $x \in D$ to $C(\theta)$. The term $\ln p(y|\theta')$ in (7) is approximated in Maybank (2006) by

$$\ln p(y|\theta') = -\frac{1}{4t} w(y, \theta')^2 + O(t^0). \tag{8}$$

It follows from (6),(7) and (8) that

$$K(\theta) = \frac{1}{4t\, l(\theta)} \int_0^{2\pi} \left(\partial^2_{\theta'\theta'} w(y, \theta')^2 \right)_{y=\tilde{x}(\phi), \theta'=\theta} a(\theta, \phi)\, d\phi, \tag{9}$$

where the $O(t^0)$ terms are omitted.

The expression (9) for $K(\theta)$ is rewritten as follows. Let x be a point in D, let $\phi \mapsto q(\theta', \phi)$ be the parameterisation of the curve $C(\theta')$, as defined in (4) and let $f(x, \theta')$ be defined by

$$f(x, \theta') = (x - q(\theta', \phi)).n(\theta', \phi). \tag{10}$$

The value of $f(x, \theta')$ is obtained by projecting $x - q(\theta', \phi)$ onto $n(\theta', \phi)$. If x is near to $q(\theta', \phi)$, then $f(x, \theta')$ is an estimate of the signed distance $\pm w(x, \theta')$. The estimate is correct to first order in $x - q(\theta', \phi)$. Let $v(\theta, \phi)$ be defined by

$$v(\theta, \phi) = \partial_{\theta'} f(x, \theta')|_{x=q(\theta,\phi), \theta=\theta'}. \tag{11}$$

The Taylor expansion of $f(x, \theta')$ about θ is

$$f(x, \theta') = f(x, \theta) + \sum_{i=1}^n v_i(\theta, \phi)(\theta'_i - \theta_i) + O_2, \tag{12}$$

where O_2 is second order or higher in the $\theta'_i - \theta_i$ and n is the dimension of the manifold T. It follows from (4),(10) and (11) that

$$\begin{aligned}
v(\theta, \phi) &= \partial_{\theta'} f(x, \theta')|_{x=q(\theta,\phi), \theta'=\theta}, \\
&= \left((\partial_{\theta'}(x - q(\theta', \phi))).n(\theta', \phi) + (x - q(\theta', \phi)).\partial_{\theta'} n(\theta', \phi) \right)_{x=q(\theta,\phi), \theta'=\theta}, \\
&= -(\partial_\theta q(\theta, \phi)).n(\theta, \phi), \\
&= -n^\top(\theta, \phi)(\partial_\theta L(\theta))z(\phi) - n(\theta, \phi).(\partial_\theta m(\theta)). \tag{13}
\end{aligned}$$

It follows from (10) that

$$\begin{aligned}
\partial^2_{\theta'\theta'} w(y, \theta')^2\big|_{y=q(\theta,\phi), \theta'=\theta} &= \partial^2_{\theta'\theta'} f(y, \theta')^2\big|_{y=q(\theta,\phi), \theta'=\theta}, \\
&= \left(2f(y, \theta')\partial^2_{\theta'\theta'} f(y, \theta') + 2\partial_{\theta'} f(y, \theta')\partial_{\theta'} f(y, \theta') \right)_{y=q(\theta,\phi), \theta'=\theta}, \\
&= 2v(\theta, \phi) \otimes v(\theta, \phi). \tag{14}
\end{aligned}$$

It follows from (9) and (14) that

$$K(\theta) = \frac{1}{2t\, l(\theta)} \int_0^{2\pi} v(\theta, \phi) \otimes v(\theta, \phi) a(\theta, \phi)\, d\phi. \tag{15}$$

7

3.3 Ellipses

It is assumed from now on that the image D is the open unit disc and that C is a circle of radius 1, centred at the origin. The curves $C(\theta)$ are ellipses and every ellipse in D can be specified as one of the $C(\theta)$. Let C have the parameterisation

$$z(\phi) = (z_1(\phi), z_2(\phi))^\top = (\cos(\phi), \sin(\phi))^\top, \qquad 0 \le \phi < 2\pi.$$

Any ellipse can be obtained by applying a suitable affine transformation to C, but the affine transformation is not unique. In order to ensure uniqueness, it is necessary to restrict the affine transformation to a subset T of $A(2)$. The subset T is a manifold, but it is not a submanifold of $A(2)$.

The manifold T is parameterised by $\theta = (\alpha, a, b, m_1, m_2)^\top$. The matrix $L(\theta)$ and the vector $m(\theta)$ in (3) are defined by

$$
\begin{aligned}
L(\theta) &= \begin{pmatrix} \cos(\alpha) & -\sin(\alpha) \\ \sin(\alpha) & \cos(\alpha) \end{pmatrix} \begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix}, \\
m(\theta) &= (m_1, m_2)^\top.
\end{aligned}
\tag{16}
$$

The ranges of α, a, b, m_1, m_2 are described below in Section 3.4.

The components α, a, b, m_1, m_2 of θ have the following geometric interpretations. The point $(m_1, m_2)^\top$ is the centre of $C(\theta)$, a is the length of the semi-major axis of $C(\theta)$, b is the length of the semi-minor axis and α is the angle between the z_1 coordinate axis and the major axis of $C(\theta)$.

Let $n(\theta, \phi) = (n_1(\theta, \phi), n_2(\theta, \phi))$ be the unit normal to $C(\theta)$ at the point $q(\theta, \phi)$. It follows from (13), (16) that the components of $v(\theta, \phi)$ are

$$
\begin{aligned}
v_1 &= -n(\theta, \phi) \begin{pmatrix} -a\sin(\alpha) & -b\cos(\alpha) \\ a\cos(\alpha) & -b\sin(\alpha) \end{pmatrix} \begin{pmatrix} \cos(\phi) \\ \sin(\phi) \end{pmatrix}, \\
v_2 &= -\cos(\phi)(n_1(\theta, \phi)\cos(\alpha) + n_2(\theta, \phi)\sin(\alpha)), \\
v_3 &= \sin(\phi)(n_1(\theta, \phi)\sin(\alpha) - n_2(\theta, \phi)\cos(\alpha)), \\
v_4 &= n_1(\theta, \phi), \\
v_5 &= n_2(\theta, \phi).
\end{aligned}
$$

As in Section 3.2, the probability measure dh on $C(\theta)$ is chosen to be proportional to the arc length on $C(\theta)$.

Let EllipticK be the complete elliptic integral of the first kind and let EllipticE be the complete elliptic integral of the second kind, as defined in Abramowitz and Stegun (1965),

$$
\begin{aligned}
\text{EllipticK}(m) &= \int_0^{\pi/2} (1 - m\sin^2(\theta))^{-1/2}\, d\theta, & m < 1, \\
\text{EllipticE}(m) &= \int_0^{\pi/2} (1 - m\sin^2(\theta))^{1/2}\, d\theta, & m \le 1.
\end{aligned}
$$

The names EllipticK, EllipticE are taken from Wolfram (2003). Let $(a, b) \mapsto c(a, b)$ be the function defined by

$$
\begin{aligned}
c(a, a) &= 0, & 0 < a, \\
c(a, b) &= (a^2 - b^2)^{-1}\left(a^2 + b^2 - 2a^2 \frac{\text{EllipticK}(1 - a^2/b^2)}{\text{EllipticE}(1 - a^2/b^2)}\right), & 0 < b < a. \tag{17}
\end{aligned}
$$

The expression (15) for the 5×5 symmetric matrix $K(\theta)$ is evaluated using Mathematica (Wolfram 2003). Nine of the entries of $K(\theta)$ are as follows,

$$
\begin{aligned}
K_{11} &= (a^2 - b^2)c(a,b)/(6t), \\
K_{22} &= (b^2(c(a,b) - 3) - 3a^2(c(a,b) - 1))/(12t(a^2 - b^2)), \\
K_{23} = K_{32} &= abc(a,b)/(6t(a^2 - b^2)), \\
K_{33} &= (a^2(3 + c(a,b)) - 3b^2(1 + c(a,b)))/(12t(a^2 - b^2)), \\
\begin{pmatrix} K_{44} & K_{45} \\ K_{54} & K_{55} \end{pmatrix} &= -\frac{c(a,b)}{4t} \begin{pmatrix} \cos(2\alpha) & \sin(2\alpha) \\ \sin(2\alpha) & -\cos(2\alpha) \end{pmatrix} + \frac{1}{4t} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.
\end{aligned}
\tag{18}
$$

The remaining 16 entries of $K(\theta)$ are zero. The matrix $K(\theta)$ has a block diagonal structure, with three blocks of sizes 1×1, 2×2 and 2×2, in order.

The entries of $K(\theta)$ are independent of m_1 and m_2. This is to be expected: any reasonable measure of the distance between two ellipses in the Euclidean plane should depend only on the positions of the ellipses relative to each other, not on their positions relative to the origin of an arbitrary coordinate system.

The circles in D are parameterised by a three dimensional submanifold of T for which a, m_1, m_2 are coordinates. The Riemannian metric induced on this submanifold by $K(\theta)$ is obtained from the limiting value of the squared length element $ds^2 = \theta^\top K(\theta)\theta$ as $b \to a$,

$$
ds^2 = (2t)^{-1}da^2 + (4t)^{-1}dm_1^2 + (4t)^{-1}dm_2^2.
$$

It follows that the entries of $K(\theta)$ are constant when θ takes values in the three dimensional parameter space for the circles in D. The parameter space for circles is thus flat, i.e. each point in it has an open neighbourhood isometric to an open set in the Euclidean space $\mathbb{R}^3$ (Gallot et al. 1990). A surprising generalisation of this result is obtained in Section 4.3: the entries of $K(\theta)$ are constant when θ takes values in the three dimensional submanifold of T which parameterizes the family of ellipses in D with a fixed orientation and a fixed eccentricity.

3.4   The parameter manifold for ellipses

The components of the parameter vector $\theta = (\alpha, a, b, m_1, m_2)$ are constrained as follows. The components a, b satisfy the usual constraints for the lengths of the semi-major axis and the semi-minor axis of an ellipse, $0 < b \le a$. The angle, α, can take any value in the range $0 \le \alpha < 2\pi$. However, if α is in the range $\pi \le \alpha < 2\pi$, then it can be replaced by $\alpha - \pi$ without changing the set of image points on the ellipse. For this reason, α is restricted to $0 \le \alpha < \pi$. The range $[0, \pi)$ of α can be regarded as a circle $S_{1/2}$ of radius $1/2$.

The vector θ is further constrained by the condition that $C(\theta)$ is entirely contained in D. If the length a of the semi-major axis is fixed, then $C(\theta)$ has its largest size at $a = b$, when $C(\theta)$ is a circle of radius a centred at $m(\theta)$. It follows that the constraint

$$
\|m(\theta)\| + a < 1,
\tag{19}
$$

ensures that $C(\theta)$ is contained in D. The constraint (19) has the advantage of simplicity, but it is stronger than necessary because it does exclude certain ellipses which are completely contained in D.

The final set of constraints on θ are associated with the blurring of $C(\theta)$. This blurring causes an $O(\sigma)$ uncertainty in the position of $C(\theta)$ where $\sigma = (2t)^{1/2}$ is the standard deviation of the errors in each component of $x - \tilde{x}$, as implied by (5). If a is $O(\sigma)$, then the uncertainty in the position of $C(\theta)$ is of the same order as the size of $C(\theta)$. It is difficult to distinguish between the blurred version of $C(\theta)$ and a point. Similarly, if b is $O(\sigma)$, then it is difficult to distinguish between the blurred version of $C(\theta)$ and a line segment. In both cases the theory underlying the derivation of $K(\theta)$ ceases to apply (Maybank 2006), and $K(\theta)$ is no longer an accurate approximation to the Fisher-Rao metric on T at θ. To avoid these difficulties, a and b are restricted such that it is not possible to approximate the blurred version of $C(\theta)$ by a point or by a line segment.

The exact definition of the restrictions on a, b arising from blurring are based on the medial axis of $C(\theta)$. The medial axis is the closure of the set of centres of those circles which are tangent to $C(\theta)$ at two distinct points. Equivalently, the medial axis is the closure of the set of points x for which the function $\phi \mapsto \|x - q(\theta, \phi)\|$ attains its minimum value $w(x, \theta)$ at two distinct values of ϕ. An alternative name for the medial axis is "focal locus". If any part of $C(\theta)$ is close to a point x on the medial axis, then the value of $p(x|\theta)$ will be, in general, the sum of contributions arising from two regions of $C(\theta)$ which are close together in D, but far apart, as measured within $C(\theta)$. The two regions each contain a point at which a circle centred at x is tangent to $C(\theta)$. Under these conditions, $K(\theta)$ ceases to be an accurate asymptotic approximation to the Fisher-Rao metric $J(\theta)$.

In the case of an ellipse with centre at the origin of coordinates and major axis coincident with the x_1 axis, the medial axis is

$$\{(x_1, 0), -a(1 - b^2/a^2) \le x_1 \le a(1 - b^2/a^2)\}.$$

The minimum distance from $C(\theta)$ to the medial axis is $a^{-1}b^2$. The values of a, b are restricted to ensure that $a^{-1}b^2 > 4\sigma = (32t)^{1/2}$. It follows that $a > (32t)^{1/2}$. Let $a_{inf} = (32t)^{1/2}$. The constraint $a^{-1}b^2 > (32t)^{1/2}$ reduces to

$$b^2 > a_{inf} a. \tag{20}$$

The full set of constraints on T is

$$T = \Big\{(\alpha, a, b, m_1, m_2), 0 \le \alpha < \pi, a_{inf} < a < 1, (a_{inf}a)^{1/2} < b < a,$$

$$\|(m_1, m_2)\| < 1 - a\}. \tag{21}$$

The inequalities in (21) involving a, b and (m_1, m_2) are strict to ensure that T is a manifold. There is no loss in imposing a strict inequality such as $a < 1$ in preference to the weaker inequality $a \le 1$, because the points of $A(2)$ omitted from T by imposing strict inequalities have a five dimensional volume equal to zero.

The canonical density $\tau(\theta)$ defined on T by the Riemannian metric $K(\theta)$ (Chavel, 1996; Gallot et al., 1990) is

$$\tau(\theta) = \sqrt{|\det(K(\theta))|}, \qquad \theta \in T. \tag{22}$$

It follows from (18) and (22) that $\tau(\theta)$ depends only on the components a, b of θ. For this reason it is convenient to write $\tau(a, b)$ in place of $\tau(\theta)$. The volume, $V(T, K)$, of T under $K(\theta)$ is given by

$$
\begin{aligned}
V(T, K) &= \int_T \tau(a, b)\, d\theta, \\
&= \pi^2 \int_{a_{inf}}^1 \int_{(a_{inf}a)^{1/2}}^a (1 - a)^2 \tau(a, b)\, db\, da.
\end{aligned}
\tag{23}
$$

The integral on the right hand side of (23) can be evaluated numerically for any given value of t.

The volume $V(T, K)$ is a fundamental measure of the complexity of the problem of detecting ellipses. If $V(T, K)$ is large, then detection is computationally expensive because a large number of candidate ellipses have to be checked. The only way of reducing the computation is to use additional information about the nature of the image or about the possible locations of ellipses within it.

4 Sampling the Parameter Manifold

One strategy for sampling a space is to begin with efficient samplings of subspaces and then build up from them to a sampling of the whole space. An example of this strategy is described by Rucklidge (1997): a sampling of the affine group $A(2)$ is built up from a sampling of the four entries of the matrix $L(\theta)$ and a sampling of the two coordinates of $m(\theta)$, where $L(\theta)$, $m(\theta)$ are as defined in (3). An examination of Conway and Sloane (1999) strongly suggests that it is not possible to build up efficient samplings of high dimensional spaces in this way. Instead, the following strategy for sampling T is adopted: choose the best known sampling of $\mathbb{R}^5$ and then map it to T. The construction of the map relies heavily on the approximation $K(\theta)$ to the Fisher-Rao metric on T.

4.1 Covering radius

Let U be a Riemannian manifold and let $\text{dist}_U(\theta, \theta')$ be the distance between points θ, θ' of U. If θ is fixed, and if $\text{dist}_U(\theta, \theta')$ is sufficiently small, then $\text{dist}_U(\theta, \theta')$ is the length of the unique arc-length parameterised geodesic in U from θ to θ'. Let G be a non-empty discrete set of points in U. Conway and Sloane (1999) define the covering radius, $r(U, G)$, of G by

$$
r(U, G) = \sup_{\theta \in U} \inf_{\theta' \in G} \text{dist}_U(\theta, \theta').
\tag{24}
$$

If θ is any point in U, then there exists an element θ' of G such that $\text{dist}_U(\theta, \theta') \le r(U, G)$. Let $B_r(\theta)$ be the subset of U defined by

$$
B_r(\theta) = \{\theta', \theta' \in U \text{ and } \text{dist}_U(\theta, \theta') < r\}.
\tag{25}
$$

It follows from (24) and (25) that $r(U, G)$ is the infimum of the numbers r for which

$$
U = \bigcup_{\theta \in G} B_r(\theta).
$$

11

The covering radius is the key parameter for measuring the effectiveness of the set G as an approximation to U. If $r(U, G)$ is small, then every point in U is close to at least one point of G.

There is some latitude in choosing a value for the covering radius of the set of sample points in T. A convenient and simple default value is to choose a value equal to 1. If $\text{dist}_T(\theta, \theta') \leq 1$, then it follows from (1) that the leading order term of the average of the log likelihood ratio $\ln(p(x, \theta)/p(x, \theta'))$ is less than or equal to $1/2$. The task is to find a set $G \subset T$ for which

$$G = \text{argmin}_{r(T, G') \leq 1} G' \mapsto |G'|. \tag{26}$$

In this application G is a finite set because the volume of T is finite under $K(\theta)$.

4.2 The lattice A_5^*

The problem of finding a G which satisfies (26) is unsolved. It remains unsolved even if T is replaced by the Euclidean space $\mathbb{R}^5$ (Conway and Sloane 1999). (In replacing T by $\mathbb{R}^5$ it is necessary to take account of the fact that $\mathbb{R}^5$ has an infinite volume. This is done by counting the number of points of G within a large ball centred at the origin of $\mathbb{R}^5$.) If the set G is restricted such that the points in G are vertices of a lattice in $\mathbb{R}^5$, then an optimal set G for $\mathbb{R}^5$ is known.

Let $v(1), \ldots, v(m)$ be a set of linearly independent vectors in $\mathbb{R}^n$, let Z be the set of integers, and let L be the set of points in $\mathbb{R}^n$ defined by

$$L = \left\{ \sum_{i=1}^{m} n(i)v(i), n(i) \in Z \right\}.$$

The set L is a lattice in $\mathbb{R}^n$. If $U = \mathbb{R}^5$, under the Euclidean metric, then the optimal lattice for sampling U is $G = A_5^*$, in the notation of Conway and Sloane (1999). The lattice A_5^* is known as Voronoi's principal lattice of the first type. It is generated as a sublattice of $\mathbb{R}^6$ by the row vectors $v(i)$, $1 \leq i \leq 5$ of the matrix

$$M' = \sqrt{\frac{72}{35}} \begin{pmatrix} 1 & -1 & 0 & 0 & 0 & 0 \\ 1 & 0 & -1 & 0 & 0 & 0 \\ 1 & 0 & 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 & -1 & 0 \\ -5/6 & 1/6 & 1/6 & 1/6 & 1/6 & 1/6 \end{pmatrix}. \tag{27}$$

The factor $\sqrt{72/35}$ is included in (27) to ensure that $r(\mathbb{R}^5, G) = 1$. The $v(i)$ span the five dimensional subspace H of $\mathbb{R}^6$ consisting of those vectors for which the sum of the coordinates is zero.

The fundamental parallelotope of A_5^* is defined by

$$\left\{ \sum_{i=1}^{5} r(i)v(i), 0 \leq r(i) < 1 \right\}.$$

The subspace H of $\mathbb{R}^6$ is covered by translates of the fundamental parallelotope, such that each translate contains a single point of the lattice. Let $\Lambda' = M'M'^{\top}$. The volume

of the fundamental parallelotope of A_5^* is $\sqrt{\det(\Lambda')}$. This suggests, but does not prove, that the number of points in an efficient sampling of T is of the order

$$n_s \equiv V(T, K)/\sqrt{\det(\Lambda')}. \tag{28}$$

where $V(T, K)$ is given by (23). A graph of n_s as a function of t is shown in Fig. 2. If $t = 10^{-3}$, then $n_s = 6492$, rounded to the nearest integer.
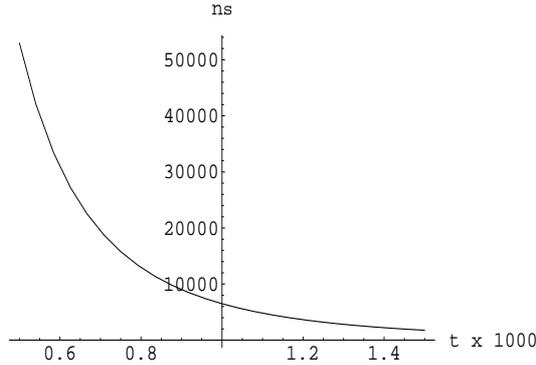

Figure 2. Graph of $n_s = V(T, K)/\sqrt{\det(\Lambda')}$ as a function of t.

The calculations in this subsection are applicable to the design of a set of Hough accumulators for ellipse detection. Each accumulator should be of the form $B_r(\theta)$. If $r = 1$, then the total number of accumulators is n_s, as defined by (28).

4.3 Definition of the sample points

The sample points in T are obtained by mapping a subset of A_5^* to T. The map is constructed by first finding a three dimensional submanifold of T on which $K(\theta)$ is constant and then sampling this submanifold using a three dimensional sublattice of A_5^*. The sampling is then extended from the submanifold to the whole of T.

The first step is to define the submanifold of T on which $K(\theta)$ is constant. Let ψ be a point of T and let $U(\psi)$ be the three dimensional submanifold of T defined by

$$U(\psi) = \{\theta, \theta \in T, \theta_1 = \psi_1, \theta_2/\theta_3 = \psi_2/\psi_3\}. \tag{29}$$

The points of $U(\psi)$ are given by $(\psi_1, a, a\psi_3/\psi_2, m_1, m_2)$. Each point corresponds to an ellipse with a fixed eccentricity ψ_2/ψ_3 and a fixed orientation ψ_1 of the major axis. The parameters for $U(\psi)$ are a, m_1, m_2. The manifold $U(\psi)$ has a Riemannian metric induced on it by the Riemannian metric $K(\theta)$ on T. The squared length element of the induced metric is

$$\frac{1}{4t}\left((1 + b^2/a^2 - c(a,b)(1 - b^2/a^2))da^2 + dm_1^2 + dm_2^2 - c(a,b)\cos(2\psi_1)(dm_1^2 - dm_2^2)\right) -$$

$$\frac{1}{2t}c(a,b)\sin(2\psi_1)dm_1dm_2, \tag{30}$$

where $b = a\psi_3/\psi_2$ and $c(a, b)$ is defined by (17). The term $c(a, b)$ is a function of the ratio a/b which is constant on $U(\psi)$. It follows from (30) that under the parameterisation

13

(a, m_1, m_2) of $U(\psi)$, the matrix for the Riemannian metric induced on $U(\psi)$ by $K(\theta)$ is constant. The manifold $U(\psi)$ is thus a flat manifold: each point of $U(\psi)$ is contained in a neighbourhood which is isometric to an open subset of the Euclidean space $\mathbb{R}^3$.

If θ, θ' are two points close together in $U(\psi)$, then in general $\mathrm{dist}_{U(\psi)}(\theta, \theta') > \mathrm{dist}_T(\theta, \theta')$. The reason for this strict inequality is that $U(\psi)$ is not a geodesic submanifold of T. The geodesic joining θ and θ' in T is in general strictly shorter than the geodesic joining θ, θ' in $U(\psi)$.

A new basis for A_5^* is obtained by applying the LatticeReduce function of Mathematica (Wolfram 2003) to the rows of the matrix M' in (27). The aim of LatticeReduce is to find a basis which consists only of short vectors. The basis of A_5^* found by LatticeReduce is given by the rows of the following matrix,

$$M = \sqrt{\frac{2}{35}} \begin{pmatrix} -5 & 1 & 1 & 1 & 1 & 1 \\ 1 & -5 & 1 & 1 & 1 & 1 \\ 1 & 1 & -5 & 1 & 1 & 1 \\ 1 & 1 & 1 & -5 & 1 & 1 \\ 1 & 1 & 1 & 1 & -5 & 1 \end{pmatrix}. \tag{31}$$

The rows of M all have length $(12/7)^{1/2}$. In contrast, one row of M' has length $(12/7)^{1/2}$ and the remaining rows have the common length $12/(35)^{1/2}$ which is larger than $(12/7)^{1/2}$. The volume of the fundamental parallelotope is independent of the choice of basis for A_5^*, $\det(MM^\top) = \det(M'M'^\top)$. The advantage of the basis of A_5^* defined by M over the basis defined by M' is that the fundamental parallelotope is closer in shape to a hypercube. The sampling algorithm described below maps a subset of A_5^* into T and in effect covers T by translates of deformed versions of the fundamental parallelotope. This covering of T is better behaved if the fundamental parallelotope is close in shape to a hypercube, and worse behaved if the fundamental parallelotope is highly skewed, i.e. very extended in one or more directions and very compressed in other directions.

A value of ψ is chosen inside T and away from the boundary of T. Let $v(i)$, $1 \le i \le 5$ be the rows of M and let $\mathrm{Tan}_\psi T$ be the tangent space to T at ψ. Vectors $u(i)$, $1 \le i \le 5$ are found in $\mathrm{Tan}_\psi T$ such that

$$u(i)^\top K(\psi) u(j) = v(i).v(j), \qquad 1 \le i, j \le 5. \tag{32}$$

The equations (32) ensure that A_5^* is mapped isometrically to $\mathrm{Tan}_\psi T$ by the linear map defined on the basis of A_5^* by $v(i) \mapsto u(i)$, $1 \le i \le 5$. Note that constraints on the $u(i)$ in addition to those of (32) are imposed in the remaining part of this subsection. Let P be the projection from $\mathbb{R}^6$ to $\mathbb{R}^5$ which removes the last coordinate, let R be a 6×6 orthogonal matrix such that $(Rv(i))_6 = 0$, $1 \le i \le 5$ and let

$$u(i) = K(\psi)^{-1/2} P R v(i), \qquad 1 \le i \le 5.$$

It can be verified that (32) holds with this definition of the $u(i)$.

Next, additional constraints are imposed on R to ensure that $u(1)$, $u(2)$, $u(3)$ are in $\mathrm{Tan}_\psi U(\psi)$. It follows from the definition (29) of $U(\psi)$ that $u(1)$, $u(2)$, $u(3)$ are in $\mathrm{Tan}_\psi U(\psi)$ if

$$\begin{aligned} u(i)_1 &= 0, & 1 \le i \le 3, \\ u(i)_2 - (\psi_1/\psi_2) u(i)_3 &= 0, & 1 \le i \le 3. \end{aligned} \tag{33}$$
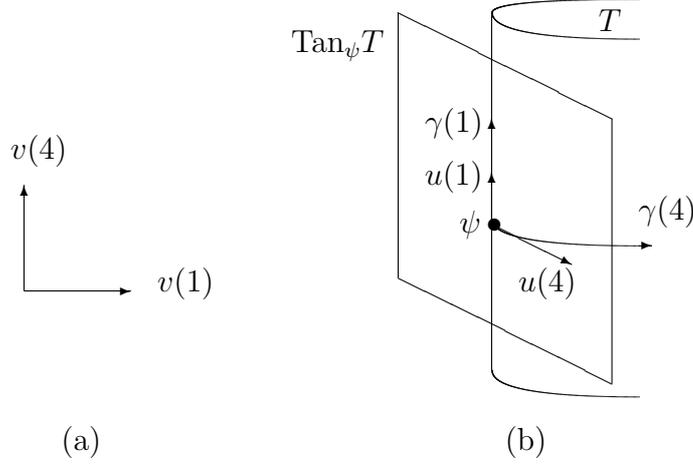
14

Figure 3. The embedding of a subset of the lattice A_5^* in T. (a) Vectors $v(1)$, $v(4)$ in $\mathrm{I\!R}^6$. (b) vectors $u(1)$, $u(4)$ in $\mathrm{Tan}_\psi T$. The vectors $u(1)$, $u(4)$ are tangent to geodesics $\gamma(1)$, $\gamma(4)$, respectively. Note that $\gamma(1)$ is tangent to $\mathrm{Tan}_\psi T$ at every point but $\gamma(4)$ curves away from $\mathrm{Tan}_\psi T$.

It follows from the block diagonal structure of $K(\psi)$, as shown in (18), that (33) holds if the following constraints on the orthogonal matrix R hold,

$$
\begin{aligned}
(Rv(i))_1 &= 0, & 1 \le i \le 3, \\
s_1(Rv(i))_2 - s_2(Rv(i))_3 &= 0, & 1 \le i \le 3,
\end{aligned}
\tag{34}
$$

where s_1, s_2 are functions of $K(\psi)^{-1/2}$ and ψ_2/ψ_3. A suitable R can always be found. The geometry underlying the choice of the vectors $u(i)$ is illustrated in Fig. 3 for $u(1)$, $u(4)$.

The manifold $U(\psi)$ is sampled at the points

$$
G(\psi) = \left\{ \theta, \theta \in U(\psi) \text{ and } \theta = \psi + \sum_{i=1}^{3} n(i)u(i), n(i) \in Z, i = 1, 2, 3 \right\}.
\tag{35}
$$

Let θ, θ' be points in $G(\psi)$. It follows that

$$
\mathrm{dist}_{U(\psi)}(\theta, \theta') = \mathrm{dist}_{\mathrm{I\!R}^6}\left(\sum_{i=1}^{3} n(i)v(i), \sum_{i=1}^{3} n'(i)'v(i) \right).
$$

As noted earlier in this subsection, $U(\psi)$ is not a geodesic submanifold of T, and in general

$$
\mathrm{dist}_{U(\psi)}(\theta, \theta') > \mathrm{dist}_T(\theta, \theta').
$$

The points of $G(\psi)$ are thus closer together in T than strictly necessary.

The parts of T outside $U(\psi)$ are sampled as follows. Let θ be a point in $G(\psi)$. The vectors $u(4)$, $u(5)$ are parallel transported from $\mathrm{Tan}_\psi T$ to $\mathrm{Tan}_\theta T$ along a curve in $U(\psi)$ from ψ to θ. Let $u(\theta, 4)$, $u(\theta, 5)$ be the resulting vectors in $\mathrm{Tan}_\theta T$. Recall that the

15

exponential map $\exp_\theta$ is defined on an open neighbourhood of the origin in $\mathrm{Tan}_\theta T$ and that if the open neighbourhood is sufficiently small, then the neighbourhood is mapped diffeomorphicly onto its image in T (Chavel, 1996; Gallot et al., 1990). The manifold T is sampled at the points

$$G = \bigcup_{\theta \in G(\psi)} \{\theta', \theta' \in T \text{ and } \theta' = \exp_\theta(n(4)u(\theta,4) + n(5)u(\theta,5)), n(i) \in Z, i = 4,5\}. \quad (36)$$

The sampling defined by (36) is most efficient in those regions near to $U(\psi)$ in which the variation in $K(\theta)$ is small.

4.4 Approximation to reduce the time needed to sample T

The computational cost of locating the points θ' in (36) is high, because it is necessary to examine many geodesics in T which begin at θ and have initial tangent vectors parallel to $n(4)u(\theta,4) + n(5)u(\theta,5)$ for $n(4)$, $n(5)$ in Z. If the norm of $n(4)u(\theta,4) + n(5)u(\theta,5)$ is too large, then the geodesic leaves T before reaching θ'. In such cases the time spent calculating the geodesic is wasted. In order to find G efficiently, it is necessary to restrict the set of values of $n(4)$, $n(5)$ for which calculation of the geodesic is attempted. Let $W(\theta)$ be the subspace of $\mathrm{Tan}_\theta T$ spanned by $u(\theta,4)$, $u(\theta,5)$. The set of values of $n(4)$, $n(5)$ is restricted by first approximating the boundary of $\exp_\theta^{-1}(T) \cap W(\theta)$ by an ellipse E and then choosing only those values of $n(4)$, $n(5)$ for which $n(4)u(\theta,4) + n(5)u(\theta,5)$ is within E.
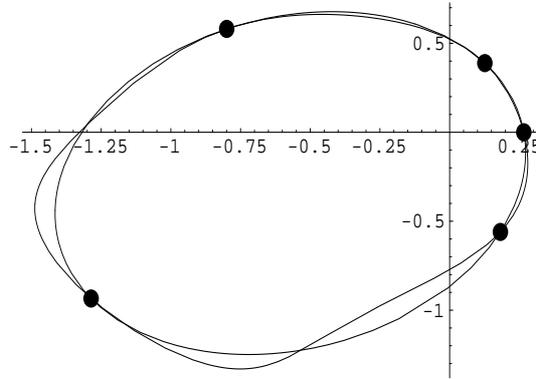
Figure 4. The points $r(i)w(i)$, $1 \le i \le 5$ are shown as dots on the unique ellipse E which passes through them. The irregular closed curve represents the boundary of $\exp_\theta^{-1}(T) \cap W(\theta)$. Note that the points $r(i)w(i)$ are closer together on that part of E near to the origin.

The following method is used to find E. Unit vectors, $w(i)$, $1 \le i \le 5$, are chosen in $W(\theta)$, such that the angle between each consecutive pair $w(i)$, $w(i+1)$ is close to $2\pi/5$. The $w(i)$ are, at least approximately, evenly spaced around the origin in $W(\theta)$. Each $w(i)$ specifies a unique geodesic in T, starting at θ and with initial tangent vector $w(i)$,

$$r \mapsto \exp_\theta(rw(i)), \qquad r \ge 0. \quad (37)$$

Let $r(i)$ be the value of r at which the geodesic (37) meets the boundary of T for the first time. If the five points $r(i)w(i)$ lie on a unique ellipse, as they do in most cases, then this

16

ellipse is chosen to be E, as illustrated in Fig. 4. The exact definition of the vectors $w(i)$ is

$$w(i) = s(i)(\cos(2\pi(i-1)/5)u(\theta,4) + \sin(2\pi(i-1)/5)u(\theta,5)), \qquad 1 \leq i \leq 5,$$

where $s(i)$ is chosen such that $w(i)^\top K(\theta)w(i) = 1$. The points in E of the form $u = n(4)u(\theta,4) + n(5)u(\theta,5)$, $n(4), n(5) \in Z$ are found. If $\theta' = \exp_\theta(u)$ is in T, then it is chosen as a sample point. If the geodesic with initial tangent vector in the direction u leaves T before reaching θ' then a sample point is chosen on the geodesic shortly before it leaves T.

If two or more of the points $r(i)w(i)$, $1 \leq i \leq 5$ are close to the origin and therefore close together, then the five points may fail to lie on an ellipse. In such cases it is necessary to choose five points more evenly spaced around the boundary of $\exp_\theta^{-1}(T) \cap W(\theta)$. Let $r(i_1)$ be the minimum of the $r(i)$, $1 \leq i \leq 5$ and let $r(i_2)$, $r(i_3)$ be the two largest values of the $r(i)$. A circle is drawn through $r(i_1)w(i_1)$, $r(i_2)w(i_2)$, $r(i_3)w(i_3)$. Then five points $\tilde{w}(i)$ are chosen equally spaced on the boundary of the circle and scalars $\tilde{r}(i)$ are found such that the points $\exp_\theta(\tilde{r}(i)\tilde{w}(i))$ are on the boundary of T. The ellipse E is chosen to pass through the five points $\tilde{r}(i)\tilde{w}(i)$, the points $u = n(4)u(\theta,4) + n(5)u(\theta,5)$, $n(4), n(5) \in Z$ are found in E and sample points θ' chosen in T as before. In the experiments there were no cases in which the points $\tilde{r}(i)\tilde{w}(i)$ failed to lie on a unique ellipse.

An example of an ellipse E containing the points $n(4)u(\theta,4) + n(5)u(\theta,5)$ is shown in Fig. 5. The axes in Fig. 5 correspond to $u(\theta,4)$ and $u(\theta,5)$. In this example, θ is chosen in $G(\psi)$ such that E contains the largest number of points $n(4)u(\theta,4) + n(5)u(\theta,5)$, $n(4), n(5) \in Z$. The true angle between the axes is $\cos^{-1}(1/5)$ but they are shown as orthogonal. The boundary of $\exp_\theta^{-1}(T)$ is shown by the small dots and the points $n(4)u(\theta,4) + n(5)u(\theta,5)$ in E are shown by the large dots.
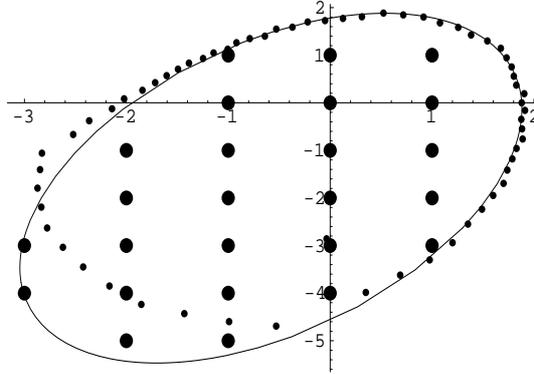

Figure 5. The boundary of $\exp_\theta^{-1}(T) \cap W(\theta)$ (small dots) and an approximating ellipse E (continuous curve). The large dots are points $n(4)u(\theta,4) + n(5)u(\theta,5)$ for certain values of $n(4)$, $n(5)$ in Z.

The number of geodesics in the above calculation is, to a first approximation, proportional to the number n_s of lattice points. Thus the number of geodesics increases rapidly as the noise level t decreases, as illustrated in Fig. 2. The length of the geodesics is proportional to $t^{-1/2}$, thus the average cost of calculating the geodesics also increases as t decreases.

4.5   Summary of the algorithm for sampling T

The algorithm for sampling T is summarised in this subsection. Recall that the basis vectors of the lattice A_5^* are chosen to be the rows $v(i)$, $1 \leq i \leq 5$ of the matrix M in (31) and $U(\psi)$ is a three dimensional submanifold of T such that ψ is in $U(\psi)$ and the matrix $K(\theta)$ is constant for θ in $U(\psi)$. It is assumed that the value of t is known.

1. Choose a point ψ in the interior of T.

2. Find an isometry $\tilde{R}$ from the Euclidean space spanned by the $v(i)$, $1 \leq i \leq 5$ to $\mathrm{Tan}_\psi T$ such that $\tilde{R}v(i)$ is in $\mathrm{Tan}_\psi U(\psi)$ for $i = 1, 2, 3$.

3. Let $u(i) = \tilde{R}v(i)$, $1 \leq i \leq 5$. Find the set $G(\psi)$ defined by

$$G(\psi) = \left\{ \theta, \theta \in T \text{ and } \theta = \psi + \sum_{i=1}^{3} n(i)u(i) \text{ with } n(i) \in Z, i = 1, 2, 3 \right\}.$$

4. For each $\theta \in G(\psi)$,

   4.1. Parallel transport $u(4)$, $u(5)$ through $U(\psi)$ from $\mathrm{Tan}_\psi T$ to $\mathrm{Tan}_\theta T$ to give vectors $u(\theta, 4)$, $u(\theta, 5)$ in $\mathrm{Tan}_\theta T$. Let $W(\theta)$ be the subspace of $\mathrm{Tan}_\theta T$ spanned by $u(\theta, 4)$, $u(\theta, 5)$.

   4.2. Approximate the boundary of $\exp_\theta^{-1}(T) \cap W(\theta)$ by an ellipse E and find all points u in E of the form $u = n(4)u(\theta, 4) + n(5)u(\theta, 5)$ with $n(4), n(5) \in Z$.

   4.3. For each point u, find the geodesic γ defined by $r \mapsto \exp_\theta(ru)$, $0 \leq r \leq 1$. If γ remains in T set $\theta' = \gamma(1)$. If γ leaves T, then choose θ' on γ shortly before it leaves T for the first time.

5. Return the union G of $G(\psi)$ and all the points θ' of T found at step 4.3.

5   Ellipse Detection

There are many different criteria for detecting an ellipse in an image. For example, an ellipse could be detected if there are a large number of edge elements which can be concatenated together to form an arc running near to or on the ellipse. The list of references for detection based on concatenations of edge elements is overwhelmingly long. Leavers (1992) is one possible starting point for an investigation of the literature. A second criterion is to detect an ellipse if the set of pixel values inside the ellipse is different in some well defined way from the set of pixel values outside the ellipse. For example, in Arias-Castro et al. (2005) an ellipse is detected if the set of values of the pixels within it has an elevated mean value. The choice of a particular criterion depends on the application, and especially on any prior knowledge about the ellipses which are to be detected.

The criterion chosen for the current work is based on a comparison of the values of the pixels just inside the ellipse with the values of the pixels just outside the ellipse. It is assumed that the values of the pixels well inside the ellipse are not relevant to the detection of the ellipse.

5.1 Criteria for detection

In this discussion of practical criteria for ellipse detection it is necessary to consider two versions of the image. The first version is the unit disk D used in the theoretical calculations carried out in Section 3. The second version is the pixel image. The noise level $\sigma = (2t)^{1/2}$ is defined in D and each parameter vector θ specifies an ellipse in D which corresponds to an ellipse $E_p(\theta)$ in the pixel image. The subscript p refers to the pixel image. It is assumed that the pixel values are scalar, i.e. the image is a grey level image.

The estimated number n_s of sample points in T, as defined by (28), increases rapidly as the noise level σ becomes small. The noise level σ is chosen such that n_s is small enough to ensure that ellipse detection can be carried out in a reasonable time. It is assumed that the resolution of the original pixel image is equal to or higher than the resolution corresponding to the chosen value of σ.

If the initial pixel image is not square, then it is reduced to a square image by deleting rows or columns as appropriate. The image is smoothed using a Gaussian mask with standard deviation σ, and then subsampled to give an image I of size $\lfloor 2\sigma^{-1} \rfloor \times \lfloor 2\sigma^{-1} \rfloor$. Let θ be a sample point in T, let c be the centre of I,

$$c = \frac{1}{2}(\lfloor 2\sigma^{-1} \rfloor, \lfloor 2\sigma^{-1} \rfloor),$$

and let r be the radius of the largest circle contained in I, $r = \lfloor 2\sigma^{-1} \rfloor$. Let $E_p(\theta)$ be the ellipse in I and corresponding to θ. Let s be an arc length parameter on $E_p(\theta)$ and let $l(E_p(\theta))$ be the arc length of $E_p(\theta)$. The ellipse $E_p(\theta)$ is obtained from the ellipse specified in D by θ by first translating the origin of coordinates to c and then scaling by r. The ellipse $E_p(\theta)$ is given in parameterised form by the following version of (3),

$$s \mapsto q_p(s) \equiv c + r(L(\theta)z(\phi(s)) + m(\theta)), \qquad 0 \le s < l(E_p(\theta)).$$

Consider the section of $E_p(\theta)$ beginning at the point with arc length parameter s and with length $E_p(\theta)/4$. The set of pixels inside $E_p(\theta)$ and close to the section is $S(1, s)$ and the set of pixels outside E_p and close to the section is $S(2, s)$. The distributions of the grey levels of the pixels in $S(1, s)$, $S(2, s)$ are modelled using Gaussian densities, and the two Gaussian densities are compared using the Kullback-Leibler distance. The ellipse $E_p(\theta)$ is detected if the distance between the two Gaussian densities is greater than a preassigned threshold for all possible values of the starting point s. The choice of threshold is discussed below in Section 5.5.

The noise level in I corresponding to the noise level σ in the unit disk D is 1 pixel. Let $n_p(s)$ be the inward normal to $E_p(\theta)$ at $q_p(s)$ and let $S(1, s)$, $S(2, s)$ be defined by

$$S(1, s) = \bigcup_{0 \le s' \le l(E_p(\theta))/4} \mathrm{Round}[\{q_p(s + s') + n_p(s + s'), q_p(s + s') + 2n_p(s + s')\}],$$

$$S(2, s) = \bigcup_{0 \le s' \le l(E_p)/4} \mathrm{Round}[\{q_p(s + s') - n_p(s + s'), q_p(s + s') - 2n_p(s + s')\}].$$

The set $S(1, s)$ is just within the boundary of $E_p(\theta)$ and the set $S(2, s)$ is just outside the boundary of $E_p(\theta)$. The definitions of $S(1, s)$ and $S(2, s)$ involve sets containing only a finite number of elements because of the presence of the 'Round' function.

Let $\mu(i, s)$, $\sigma(i, s)$ be the mean and standard deviation of the grey levels of the pixels in $S(i, s)$ for $i = 1, 2$, and let $g \mapsto f(i, s, g)$ be the Gaussian density for the grey level g with expected value $\mu(i, s)$ and standard deviation $\sigma(i, s)$, $i = 1, 2$. The difference between the two sets of grey levels derived from the pixels in $S(1, s)$ and $S(2, s)$ is measured by the Kullback-Leibler distance $D_s(2||1)$ between $f(1, s, g)$ and $f(2, s, g)$,

$$
\begin{aligned}
D_s(2||1) &= \int_{-\infty}^{\infty} \ln\left(\frac{f(2, s, g)}{f(1, s, g)}\right) f(2, s, g)\, dg, \\
&= \ln\left(\frac{\sigma(1, s)}{\sigma(2, s)}\right) + \frac{\sigma(2, s)^2}{2\sigma(1, s)^2} - \frac{1}{2} + \frac{1}{2}\left(\frac{\mu(2, s) - \mu(1, s)}{\sigma(1, s)}\right)^2.
\end{aligned}
\tag{38}
$$

The motivating idea is that on passing from the inside of $E_p(\theta)$ to the outside the density function for the grey levels changes from $f(1, s, g)$ to $f(2, s, g)$. The Kullback-Leibler distance $D_s(2||1)$ measures the penalty in assuming the density function for the grey levels continues to be $f(1, s, g)$ when in fact it has changed to $f(2, s, g)$. The distance $D_s(2||1)$ is chosen in preference to $D_s(1||2)$ because in many applications the interior of $E_p(\theta)$ is relatively uniform, whereas the exterior is strongly patterned, depending on the appearances of those objects partially occluded by $E_p(\theta)$. In such applications, $D_s(2||1)$ tends to be large. It is noted that $D_s(2||1)$ is invariant under a uniform scaling of the image grey levels.

Let $\zeta(E_p(\theta))$ be defined by

$$
\zeta(E_p(\theta)) = \min_{0 \leq s \leq l(E_p(\theta))} D_s(2||1).
\tag{39}
$$

The min function is used in (39) rather than inf because the set of distinct values of $D_s(2||1)$ is finite. If $D_s(2||1)$ is small for some value of s, then it is likely that the grey levels of the pixels in $S(1, s)$ and $S(2, s)$ are similar and that the boundary of $E_p(\theta)$ has significant gaps along the arc beginning at s and with length $l(E_p(\theta))/4$. The ellipse $E_p(\theta)$ is detected if all the $D_s(2||1)$ are sufficiently large, or equivalently, $\zeta(E_p(\theta)$ is greater than a pre-assigned threshold. The ellipse with the largest value of $\zeta(E_p(\theta))$ is, by definition, the most salient ellipse.

This criterion for detecting $E_p(\theta)$ fails if the pixels just inside $E_p(\theta)$ have a constant or near constant grey level. In such cases $\sigma(1, s)$ is small, $D_s(2||1)$ is large and $E_p(\theta)$ is detected even though the differences between the grey levels in $S(1, s)$ and $S(2, s)$ may be invisible to the eye. These cases arise in natural images when the interior of $E_p(\theta)$ is so bright that the pixel values are at the top of the allowable range. The standard deviation $\sigma(1, s)$ is small because of the limitations of the imaging device, not because the original object corresponding to the ellipse is uniform in appearance. To avoid such cases, the difference r between the maximum grey level in I and the minimum grey level in I is calculated and $\sigma(1, s)$, $\sigma(2, s)$ are replaced by $\max\{r/25, \sigma(1, s)\}$, $\max\{r/25, \sigma(2, s)\}$ respectively, when calculating $D_s(2||1)$. The invariance of $D_s(2||1)$ under uniform scaling of the grey levels is preserved by this replacement.

5.2 Effect of the Gaussian assumption

A short calculation shows that the method described in the preceding subsection is sub-optimal if the probability density function for the grey levels of the pixels in $S(2, s)$ is not

Gaussian. To see this, let the probability density functions for the grey levels of the pixels in $S(s,1)$, $S(s,2)$ be $f(1,s,g)$, $h(2,s,g)$, respectively. It is not assumed that $h(2,s,g)$ is Gaussian, but $h(2,s,g)$ has the same expected value and the same standard deviation as $f(2,s,g)$. Let $\tilde{D}_s(2\|1)$ be the correct Kullback-Leibler distance between the distributions of pixel values inside and outside the ellipse.

$$\tilde{D}_s(2\|1) = \int_{-\infty}^{\infty} \ln\left(\frac{h(2,s,g)}{f(1,s,g)}\right) h(2,s,g)\, dg. \tag{40}$$

It follows from (38) and (40) that

$$\begin{aligned}
\tilde{D}_s(2\|1) - D_s(2\|1) &= \int_{-\infty}^{\infty} (\ln h(2,s,g)) h(2,s,g)\, dg - \int_{-\infty}^{\infty} (\ln f(2,s,g)) f(2,s,g)\, dg \\
&\quad + \int_{-\infty}^{\infty} (\ln f(1,s,g))(f(2,s,g) - h(2,s,g))\, dg. \tag{41}
\end{aligned}$$

The first term on the right hand side of (41) is the negative of the differential entropy of $h(2,s,g)$. The second term is the differential entropy of $f(2,s,g)$. The third term is zero because $f(1,s,g)$ is Gaussian and $f(2,s,g)$, $h(2,s,g)$ have the same expected value and the same standard deviation. The right hand side of (41) is thus non-negative because the Gaussian density has the largest differential entropy for a given standard deviation.

The result $\tilde{D}_s(2\|1) \geq D_s(2\|1)$ shows that ellipse detection based on $f(1,s,g)$, $f(2,s,g)$ does not produce any additional false alarms. If $D_s(2\|1)$ is above the threshold for detection, then $\tilde{D}_s(2\|1)$ is also above the threshold for detection. However, if $D_s(2\|1)$ is below the threshold and $\tilde{D}_s(2\|1)$ is above the threshold, then the result is a false negative. The ellipse should be detected but it is missed.
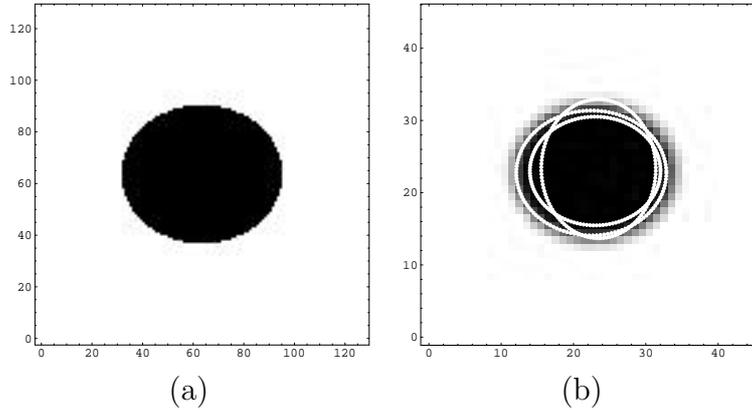

(a)   (b)

Figure 7. Ellipse detection in a synthetic image.

5.3   Experimental results

The same set G of sample points in T was used in all the experiments reported in this subsection. The value of ψ was

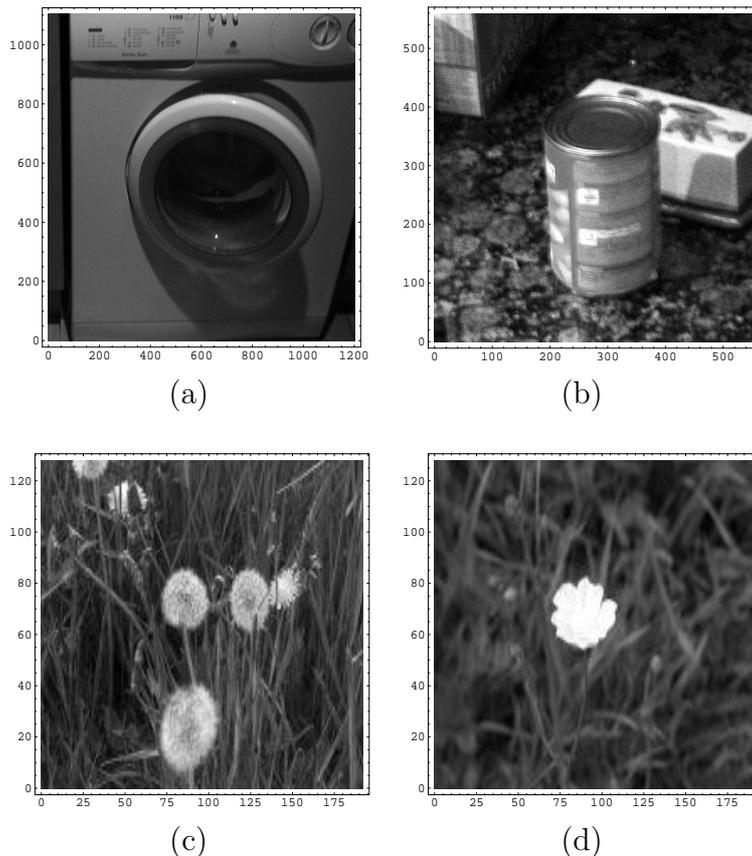$$\psi = (\pi/2, (1 - a_{inf})/2, 5(1 - a_{inf})/12, 0, 0).$$

Figure 8. Original images. (a), (b) courtesy of M.J.L. Maybank; (c), (d) are imk01184.jpg and imk01149.jpg respectively, from the image data base made available by van Hateren and van der Schaaf (1998).

The noise level σ was appropriate for an image of size 45×45, namely $\sigma = 2/45$. The size of G was $|G| = 6525$ which is close to the theoretical value $n_s = 6492$ obtained in Section 4.2. The low resolution images, of size 45×45, used for ellipse detection, were obtained from the original images using Gaussian smoothing followed by subsampling.

The ellipse detection algorithm was implemented in Mathematica, version 5, and run on a Pentium 4 PC with speed 2.80 GHz and 504 MB of RAM. The time taken to search a 45×45 grey level image for ellipses was 2 min. and 53 sec.

5.3.1. Synthetic image. Fig. 7(a) is a binary synthetic image of an ellipse. Fig. 7(b) shows the three most salient ellipses, i.e. the three ellipses with the highest values of $\zeta(E_p(\theta))$, found in the associated low resolution image of size 45×45. The values of $\zeta(E_p(\theta))$ are 26.33, 23.79 and 18.14. The closest fitting ellipse, by eye, has the highest value of $\zeta(E_p(\theta))$.

5.3.2. Natural images. The original images are shown in Fig. 8 and the low resolution images are shown in Fig. 9 with the most salient ellipses superposed on the appropriate image. In Fig. 9(a) the most salient ellipse, with $\zeta(E_p(\theta)) = 3.29$, is located on the door of the washing machine. There is also a false detection, with $\zeta(E_p(\theta)) = 3.16$, in which the ellipse is close to three strong contours in the image. The values of $\zeta(E_p(\theta))$ in Fig. 9(b)

22

are all low, consistent with the observation that the image does not contain a well defined ellipse. In Figs. 9(c) and 9(d) the flowers are detected because at this low resolution the boundaries of the flowers approximate to ellipses. The flower at the centre of Fig. 9(c) is not detected because it is slightly too small.
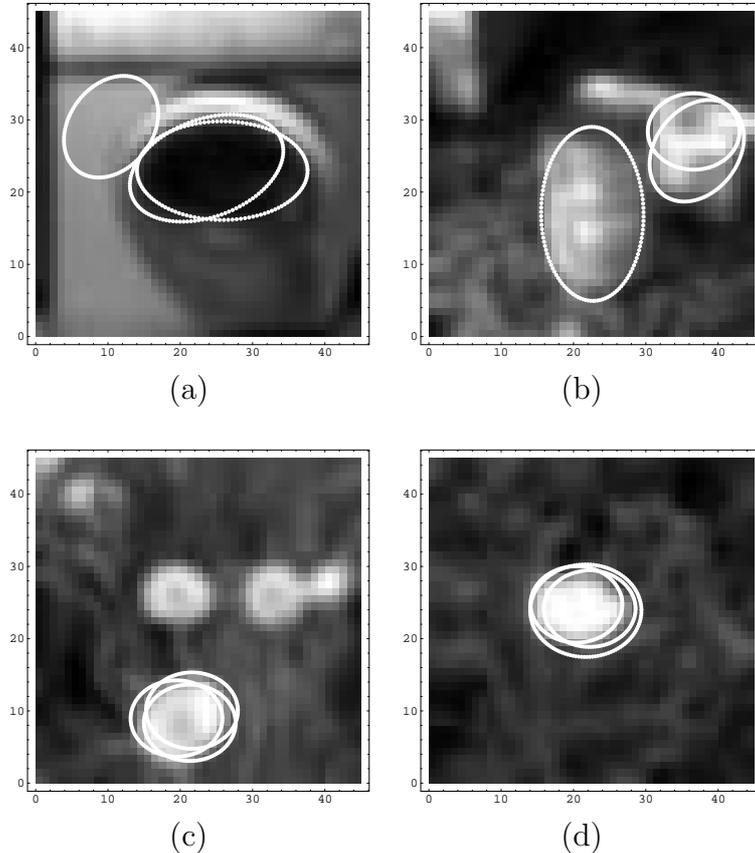

Figure 9. Downsampled images with the most salient ellipses superposed. The values of $\zeta(E_p(\theta))$ for the ellipses shown are (a) 3.29, 3.16 (false detection), 3.16; (b) 1.43, 1.34, 1.20 (on can); (c) 4.91, 2.49, 2.35; (d) 6.17, 3.65, 1.77.

5.4 Multiresolution ellipse detection

The reliability of ellipse detection can be increased using a multiresolution approach. Let I be a given image, let I_1 be a low resolution version of I, of size $m_1 \times m_1$ and let I_2 be a higher resolution version of I, of size $m_2 \times m_2$, $m_2 > m_1$. The associated measures of the noise level, after scaling I_1 and I_2 to have width 2, are $\sigma_1 = 2/m_1$ and $\sigma_2 = 2/m_2$. Let t_1, t_2 be defined by $t_1 = \sigma_1^2/2$, $t_2 = \sigma_2^2/2$ and let ξ be an ellipse detected in I_1.

The parameter manifold T is given the metric $\theta \mapsto K(t_1, \theta)$ obtained by setting $t = t_1$ in (15). Let $B_1(t_1, \xi)$ be the unit ball in T centred at ξ. The unit ball $B_1(t_1, \xi)$ carries a second metric defined by $\theta \mapsto K(t_2, \theta)$. The two metrics differ by a scale factor,

$$K(t_1, \theta) = (t_2/t_1)K(t_2, \theta), \qquad \theta \in T.$$

It is required to sample $B_1(t_1, \xi)$ such that each point of $B_1(t_1, \xi)$ is within a distance 1 of a sample point, where the distance is measured using the metric $\theta \mapsto K(t_2, \theta)$. Each

sample point defines an ellipse close to ξ in the high resolution image. The sample points are first defined in $\mathbb{R}^5$ and then mapped to T.

In order to simplify the calculations, each metric $\theta \mapsto K(t_i, \theta)$ on $B_1(t_1, \xi)$ is replaced by the appropriate flat metric $\theta \mapsto K(t_i, \xi)$, $i = 1, 2$. Let $r = \sigma_1 \sigma_2^{-1} + 1$. The value of r is chosen to be $\sigma_1/\sigma_2 + 1$ rather than σ_1/σ_2 in order to ensure that all points of the ball in $\mathbb{R}^5$ of radius σ_1/σ_2 are within a distance 1 of a sample point. Sample points x are chosen in the ball in $\mathbb{R}^5$ of radius r and centred at the origin. Each x is a vertex of a lattice of type A_5^* with covering radius equal to 1. The points x are mapped to T by

$$x \mapsto \xi + K(t_2, \xi)^{-1/2} x.$$

Note that the squared distance from ξ to $\xi + K(t_2, \xi)^{-1/2} x$ in the metric on T defined by $K(t_1, \xi)$ is approximately

$$x^\top (K(t_2, \xi)^{-1/2})^\top K(t_1, \xi) K(t_2, \xi)^{-1/2} x = (\sigma_2/\sigma_1)^2 \|x\|^2 \le (1 + \sigma_2/\sigma_1)^2.$$

Some results of multiresolution ellipse detection are shown in Fig. 10. Fig. 10(a) shows the lowest resolution image with the four most salient ellipses superposed on it. Figs. 10(b), 10(c), 10(d) show the ellipses detected at successively higher resolutions. An ellipse, θ, is shown if $\zeta(E_p(\theta)) \ge 2$ or if the ellipse is the most salient in its image. The values of $\zeta(E_p(\theta))$ for the ellipses shown are (a) 3.29, 3.16, 3.16, 2.28; (b) 2.19; (c) 1.23; (d) 0.41. The sizes of the four images are 45×45, 67×67, 99×99, 147×147. The number of ellipses sampled from T is 852 for Fig. 10(b), 213 for Fig. 10(c) and 213 for Fig 10(d).

The above multiresolution algorithm was implemented in Mathematica, version 5, and run on a Pentium 4 PC with speed 2.80 GHz and 504 MB of RAM. The time taken to generate the results shown in Fig. 10 was 2 min. and 28 sec.

5.5 Choice of threshold

It is apparent from the experimental results shown in Sections 5.3 and 5.4 that the choice of threshold for ellipse detection depends on the class of images in which the ellipses are sought. In the low resolution synthetic image shown in Fig. 7, certain values of $\zeta(E(\theta))$ can be very high, of the order of 25, because the sets of pixel values used to calculate $\zeta(E_p(\theta))$ can have very small standard deviations. In the low resolution natural images shown in Fig. 9, the local variations in grey level are much larger and a suitable threshold value for ellipse detection is 2.

In the high resolution natural images shown in Fig. 10 a new phenomenon appears. The value of $\max_{\theta \in G} \zeta(E_p(\theta))$ falls as the resolution increases. This is because the boundary of the most salient ellipse has a complicated structure. At low resolutions, the blurred version of this structure is close enough to the model for image grey levels described in Section 5.1 to give a value of $\zeta(E_p(\theta))$ greater than 2. As the resolution increases, the details of the structure appear at the ellipse boundary, the model in Section 5.1 no longer applies with the same accuracy and lower values for $\zeta(E_p(\theta))$ are obtained.

6 Conclusion

A new method for detecting ellipses in two dimensional images has been described. It relies on an efficient sampling of the five dimensional parameter space T for the ellipses in the
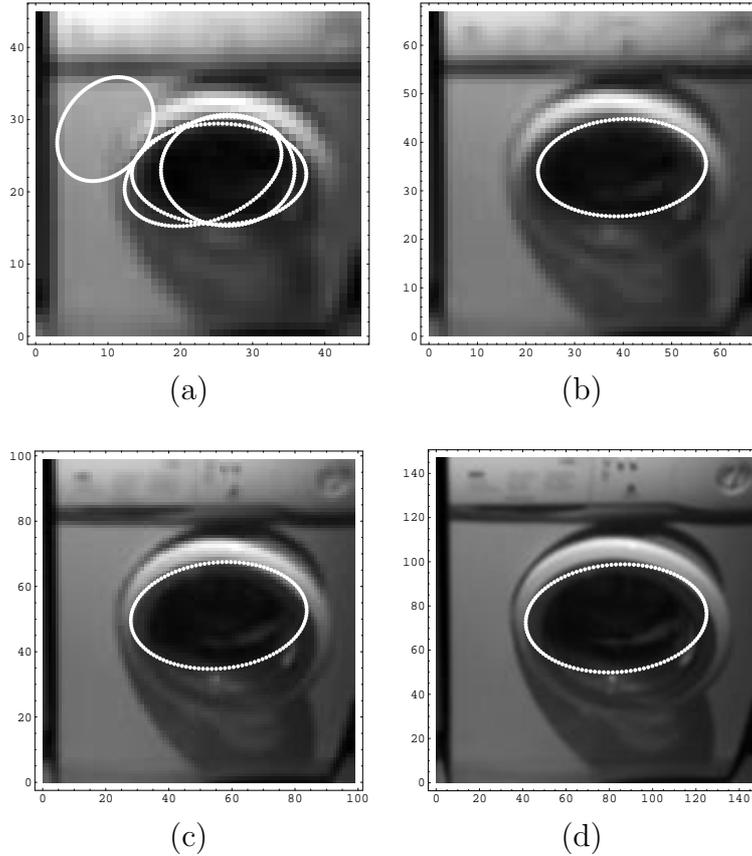
Figure 10. Ellipses detected at different resolutions.

image. The distances between nearby sample points are chosen using an approximation to a Fisher-Rao metric on T. Each sample point is checked in turn to see if the presence of the associated ellipse is supported by the image pixel values. The major advantages of this approach are:

i) It is very resistant to the effects of outliers, i.e. image structures which are not related to the ellipse which is being checked.

ii) The number of sample points chosen from T is a measure of the complexity of the task of ellipse detection. This number is proportional to the volume of T under the approximation to the Fisher-Rao metric.

The criterion for detecting an ellipse is based on the Kullback-Leibler distance between the estimated distribution of the grey levels of pixels just inside the ellipse and the estimated distribution of the grey levels of pixels just outside the ellipse. Experiments show that ellipses can be detected successfully in a variety of images. The drawbacks of this approach to ellipse detection are

i) The large number of sample points required at low noise levels.

ii) The necessity of re-sampling T if the noise level $\sigma = (2t)^{1/2}$ is changed.

25

iii) The lack of theoretical arguments to support a particular choice of detection threshold.

In the case of *(i)*, the number of sample points can be reduced by using a multiresolution algorithm, as described in Section 5.4. In the case of *(iii)*, a theoretical estimate of the detection threshold may be possible if a suitable class of random images can be defined. The threshold should be high enough to prevent the detection of ellipses in a large proportion of the random images.

There is a wide range of structure detection problems that can in principle be solved by first approximating the Fisher-Rao metric and then using the approximating metric to guide the sampling of a parameter space for the structures in question. Detailed calculations are given in Maybank (2004) for lines and in Maybank (2005) for projective transformations of the line. These methods provide a firm statistical foundation for structure detection and make it possible to compare the complexity of different detection problems using the volumes of the different parameter spaces. If the volume of a parameter space is large then the associated detection problem is in practice unsolvable unless there is further information about the image or about the structures being sought. If the volume of the parameter space is small and if there are well defined criteria for deciding if the structure in question is present, then the detection problem for that particular type of structure is solved, both theoretically and practically.

In current work there is a division between the probability model used to obtain a Fisher-Rao metric on T and the probability model used to assess the evidence supporting the presence of a given ellipse θ. The probability model used to obtain the Fisher-Rao metric is based on two maximum entropy assumptions: the true measurement is distributed uniformly on the ellipse, and the noise affecting the measurements is Gaussian. These assumptions have the advantage of simplicity, but they may preclude the full use of all the information in the image. In principle, it should be possible to use a single probability model to obtain a Fisher-Rao metric on T and to assess the evidence supporting the presence of a given ellipse.

References

Abramowitz, M. and Stegun, I.A. (eds.) 1965. *Handbook of Mathematical Functions with formulas, graphs, and mathematical tables.* Dover.

Agrawal, M. and Davis, L.S. 2003. Camera calibration using spheres: a semi-definite programming approach. In *Proc. 9th International Conference on Computer Vision,* Nice, vol. 2, pp. 782-789.

Aguado, A.S., Montiel, M.E. and Nixon, M.S. 1995. Ellipse extraction via gradient direction in the Hough transform. In *Proc. 5th Int. Conf. on Image Processing and its Applications*, Edinburgh, UK, pp. 375-378, IEE.

Amari, S.-I. 1985. *Differential-Geometrical Methods in Statistics.* Lecture Notes in Computer Science, 28. Springer.

Arias-Castro, E., Donoho, D. and Huo, X. 2005. Near-optimal detection of geometric objects by fast multiscale methods. *IEEE Transactions on Information Theory*, 51(7): 2402-2425.

Chavel, I. 1996. *Riemannian Geometry: a modern introduction.* Cambridge Tracts in Mathematics, vol. 108, CUP.

Conway, J.H. and Sloane, N.J.A. 1999. *Sphere Packings, Lattices and Groups.* Grundlehren der mathematischen Wissenschaften, vol. 290. Springer.

Cover, T.M. and Thomas, J.A. 1991. *Elements of Information Theory.* John Wiley and Sons.

Desolneux, A., Moisan, L. and Morel, J.-M. 2000. Meaningful alignments. *International Journal of Computer Vision*, 40(1): 7-23.

Desolneux, A., Moisan, L. and Morel, J.-M. 2003. Maximal meaningful events and applications to image analysis. *The Annals of Statistics*, 31(6):1822-1851.

Fitzgibbon, A., Pilu, M., and Fisher, R.B. 1999. Direct least squares fitting of ellipses. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(5): 476-480.

Gallot, S., Hulin, D. and LaFontaine, J. 1990. *Riemannian Geometry.* 2nd edition, Universitext, Springer.

Hartley, R. and Zisserman, A. 2003. *Multiple View Geometry in Computer Vision.* CUP.

van Hateren, J.H. and van der Schaaf, A. 1998. Independent component filters of natural images compared with simple cells in primary visual cortex. *Proceedings of the Royal Society of London*, Series B, 265:359-366.

Ji, Q. and Haralick, R.M. 1999. A statistically efficient method for ellipse detection. *IEEE International Conference on Image Processing, ICIP99*, vol. 2, pp. 730-734.

Kanatani, K. 1994. Statistical bias of conic fitting and renormalisation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(3): 320-326.

Kanatani, K. 1996. *Statistical Computation for Geometric Optimization.* Elsevier.

Kanatani, K. and Ohta, N. 2003. Automatic detection of circular objects by ellipse growing. *International Journal of Image and Graphics*, 4(1): 35-50.

Kwolek, B. 2004. Stereo-vision based head tracking using color and ellipse fitting in a particle filter. In Pajdla, T. and Matas, J. (eds.) 2004. *Computer Vision - ECCV 2004*, vol. 4, pp. 192-204. Lecture Notes in Computer Science, vol. 3024, Springer.

Lazebnik, S., Schmid, C. and Ponce, J. 2004. Semi-local affine parts for object recognition. In Hoppe, A., Barman, S. and Ellis, T. (eds.) *British Machine Vision Conference 2004*, vol. 2, pp. 959-968, BMVA.

Leavers, V.F. 1992. *Shape Detection in Computer Vision Using the Hough Transform.* Springer Verlag.

Leedan, Y. and Meer, P. Estimation with bilinear constraints in computer vision. *Sixth International Conference on Computer Vision*, ICCV'98, Bombay, India, pp. 733-738.

Lutton, E. and Martinez, P. 1994. A genetic algorithm for the detection of 2D geometric primitives in images. *Proc. 12th International Conference on Pattern Recognition*, vol. 1, pp. 526-529.

Maio, D. and Maltoni, D. 1998. Fast face location in complex backgrounds. In Wechsler, H., Phillips, P.J., Bruce, V., Soulié, F. and Huang, T.S. (eds.) *Face Recognition from Theory to Applications*, NATO LSI Series, pp. 568-577, Springer-Verlag.

Maybank, S.J. 2003. Fisher information and model selection for projective transformations of the line. *Proceedings of the Royal Society of London*, Series A, 459: 1-21.

Maybank, S.J. 2004. Detection of image structures using Fisher information and the Rao metric. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(12): 1579-1589.

Maybank, S.J. 2005. The Fisher-Rao metric for projective transformations of the line. *International Journal of Computer Vision*, 63(3): 191-206.

Maybank, S.J. 2006. Application of the Fisher-Rao metric to structure detection. To appear in *Journal of Mathematical Imaging and Vision.*

Rosin, P.L. 1996. Assessing error of fit functions for ellipses. *Graphical Models and Image Processing*, 58(5): 494-502.

Rucklidge, W.J. 1997. Efficiently locating objects using the Hausdorff distance. *International Journal of Computer Vision*, 24(3): 251-270.

Scaggiante, A., Frezza, R. and Zampato, M. 1999. Identifying and tracking ellipses: a technique based on elliptical deformable templates. In *Proc. 10th International Conference on Image Analysis and Processing*, pp. 582-587.

Taubin, G. 1991. Estimation of planar curves, surfaces, and non planar space curves defined by implicit equations with applications to edge and range image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(11):1115-1138.

Torr, P.H.S. and Fitzgibbon, A.W. 2004. Invariant fitting of two view geometry. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(5): 648-650.

Wang, J.G., Sung, E. and Venkateswarlu, R. 2003. Eye gaze estimation from a single image of one eye. In *Proc. 9th International Conference on Computer Vision*, Nice, vol. 1, pp. 136-143.

Wolfram, S. *The Mathematica Book*. 5th Edition, Wolfram Media, 2003.

Xie, Y. and Ji, Q. 2002. A new efficient ellipse detection method. *Proc. International Conference on Pattern Recognition*, vol. 2, pp. 957-960.

Xu, L. and Oja, E. 1993. Randomized Hough transform (RHT): basic mechanisms, algorithms and computational complexities. *Computer Vision, Graphics, and Image Processing: Image Understanding*, 57(2): 131-154.

Yao, J., Kharma, N. and Grogono, P. 2004. Fast, robust GA-based ellipse detection. In *Proc. 17th International Conference on Pattern Recognition, ICPR 2004*, vol. 2, pp. 859-862.

Younes, L. 1998. Computable elastic distances between shapes. *SIAM Journal on Applied Mathematics*, 58(2): 565-586.

Zhang, Z. 1997. Parameter estimation techniques: a tutorial with application to conic fitting. *Image and Vision Computing Journal*, 15(1): 59-76.