

Ellipse Detection and Matching with Uncertainty

Tim Ellis, Ahmed Abbood, Beatrice Brillault
Machine Vision Group,
Centre for Information Engineering, City University,
London, EC1V 0HB, UK (T.J.Ellis@uk.ac.city)

Abstract

This paper considers the fitting of ellipses to edge-based pixel data in digital images. It infers the ellipses as projected circles, viewed obliquely, and uses this information to constrain the position, viewpoint and scale of model objects projected onto the image. A central theme of the work is the explicit use of uncertainty, which reflects observation errors in the image data, and which is propagated through the model matching stages, providing a consistent and robust representation of errors.

1 Introduction

Model-based vision attempts to match image primitives (e.g. straight lines, corners, arcs) to 3D object representations, typically based on geometric models. One of the major problems in such matching lies in estimating the viewpoint of the image sensor with respect to the scene. When available, initial estimates of viewpoint are used to provide a $3D \Rightarrow 2D$ transformation of the model, followed by verification of remaining image features. Alternatively, the image primitives are used to infer 3D structures in the data, which are directly matched with the model data [1, 2].

Simple image primitives tend to provide a poor source of data from which to estimate the viewpoint. Much recent work (e.g. [3]) has used the detection of perceptual groups (e.g. parallel lines, which are broadly insensitive to viewpoint) in the image to enrich the data and constrain the range of viewpoints from which to initiate a match. In addition, this approach has been combined with a more realistic approach to matching, which employs rigorous models of uncertainty applied to the measurement data, and propagates these uncertainties through the matching process [1, 2, 7].

There has been considerable interest in recent years in the detection of ellipses in image data [4, 5]. Ellipses are commonly interpreted as oblique projections of circular features in the image [6, 8]. Such features are typically rare in scenes, and hence are excellent cues if they can be detected reliably [8]. A further advantage of circular features is that they constrain the viewpoint to a greater degree than simpler line features or groups of features (e.g. parallel lines). As such they can provide a more reliable and robust initial estimate of viewpoint.

As with parallel lines, ellipses can also be grouped in order to infer more complex underlying structures. West and Rosin [6] grouped ellipses inferred

to belong to planar surfaces (containing multiple circular features), cylindrical/conical objects (axis of revolution) and circular concentric features. Their method used the Hough paradigm to identify ellipses (projected circles) associated with these groupings.

In this paper we consider two aspects of the wider task of matching image to model primitives. Firstly, we extract ellipses from the image data and use them (singly or in groups) to estimate viewpoint and infer structures in the scene. Secondly, we match these structures to model structures. The feature extraction process, based on the Kalman filter, calculates the uncertainty associated with the ellipse fitting. This uncertainty is propagated throughout the model matching stage, in order to relate the uncertainty in the feature extraction to the location and recognition of the model object.

Initial results have been applied to suitable demonstration images which indicate the performance of the ellipse detection algorithm on a variety of projected circles. In addition, examples of ellipse grouping are also presented.

2 Ellipse Fitting

Directly fitting ellipses to pixel edge data is computationally expensive and prone to gross errors in all but the simplest images. A more practical alternative exploits connectivity of edges in the scene and initiates fitting on connected edges. Further efficiencies can be obtained by reducing connected edges to polygonal approximations, utilising only the line end-points for initial fits.

We treat ellipse fitting as a three stage process. Following application of a suitable edge detector (Marr-Hildreth, Canny, Shen [10]), 8-connected line segments are extracted from the image and straight line and arc data are generated [4]. Ellipses are initially fitted to the detected arc segments using a least-mean-square (LMS) error fitting to a general conic function, which maintains an estimate (uncertainty) of the fit through the covariance matrix. Initially fitting to the line data results in a more efficient calculation.

These initial fits may be improved by extending the arcs using existing edge connectivity information and more global information of unconnected line segments. In order to allow updating of the LMS algorithm, we use a recursive form, the Kalman filter. Initialising this algorithm with the parameters from the initial LMS fit (and in particular, the covariance matrix) avoids numerical instabilities in the Kalman filter. In addition, the Kalman filter is updated with the original pixel edge data points, rather than the polygonal end-points, in order to improve the accuracy of fit.

We describe the ellipse by the normal parametric form for a conic section:

$$f(x_i, y_i) = ax^2 + bxy + cy^2 + dx + ey + f = 0 \quad (1)$$

selecting the simple normalisation of $f=1$, which merely precludes representation of ellipses which pass through the image origin (i.e. (0,0) co-ordinate). Hence, the ellipse is described by the vector $x = (a \ b \ c \ d \ e \ 1)$ and the updating of the parameter vector and its covariance are given by

$$x_i = x_{i-1} - (z_i - h_i^t x_{i-1}) \frac{(S_{i-1} h_i)}{\sigma^2 + h_i^t S_{i-1} h_i} \quad (2)$$




Figure 1: Criteria for matching global line data to ellipse.

$$S_i = S_{i-1} - \frac{(S_{i-1}h_i)(S_{i-1}h_i)^t}{\sigma^2 + h_i^2 S_{i-1}h_i} \quad (3)$$

where

$$z = -(ax^2 + bxy + cy^2 + dx + ey + 1) \quad (4)$$

$$h^t = (x^2 \ xy \ y^2 \ x \ y \ 1)^t \quad (5)$$

$$\sigma^2 = 4n^2((ax + by + d)^2 + (bx + cy + e)^2) \quad (6)$$

and n^2 is an estimate of the spatial noise, for which we currently use a value of 0.1 (see [5]).


The iterative Kalman filter is initialised (x_0, S_0) with values calculated from the LMS method.

In extending the arcs, we use a Mahalanobis distance (MD) measure to select candidate line segments of suitable orientation and location (with respect to the current ellipse fit). The MD test weights the difference between the current fit and a candidate feature (line segment) by the current estimate of the uncertainty, and hence provides a metric for selecting features with a given statistical sampling. As shown in figure 1, we use a mid-point representation for the line segments, estimating the tangent to the data from the direction of the line. To minimise errors, the position of the line is determined by the intersection of the projection of the tangent normal with the pixel data. On accepting the segment, the Kalman filter is again updated with the pixel edge data associated with the line segment.

The MD test is given by :

$$\frac{f^2(x_i, y_i)}{\sigma_f^2} + \frac{g_{x,y}^2(u, v)}{\sigma_g^2} < t \quad (7)$$

where $f(x_i, y_i)$ is defined as in equation 1, σ_f^2 is the variance of f with respect to the euclidean co-ordinates (x_i, y_i) and the ellipse parameters, σ_g^2 is the variance



a - major axis
 b - minor axis
 c - centre
 θ - orientation
 τ - tilt angle

Figure 2: Derived parameters of ellipse.

of g with respect to the euclidean co-ordinates, the ellipse parameters and the tangential directions (u,v) , and

$$g_{x,y}^2(u,v) = 2axu + b(yu + xv) + 2cyv + du + ev \quad (8)$$

We select a value for the threshold t (from a χ^2 distribution with two degrees of freedom) in order to provide a suitable trade-off between accepting a high proportion of good features, and rejecting poor ones. In noisier images, the threshold can be reduced to make the test more selective.

A disadvantage of only fitting connected edge data is that connectivity of edges is not assured, due to noise and other interfering effects. Hence, the final stage of our fitting process searches for global (unconnected) line segments in the image that will improve the ellipse fit using the same Mahalanobis distance test given above.

3 Model Fitting

For simplicity, we assume that the projection is orthographic. Using parameters derived from the equation of the ellipse, we can infer the radius of the projected circle (the major axis), the angle of the oblique viewpoint (cosine (minor/major), the tilt), and the orientation of the plane (with respect to the viewpoint) in which the corresponding 3D circular feature lies (see figure 2).

Following the detection of multiple ellipses, grouping of the ellipses can proceed. We avoid the use of the Hough transform for grouping in favour of the Mahalanobis distance/Kalman filter combination.

The results of ellipse detection are processed to infer 3D structures in the image and the the camera viewpoint. Initially the ellipses are grouped on the basis of similar tilt and orientation. In 3D, such groupings may be attributed to co-planar or parallel-planar sets. The MD test is performed on $P(\tau, \theta)$, and is equal to $P^t C_P^{-1} P < t$ (with two degrees of freedom), where C_P is the covariance matrix associated with P . It selects appropriate ellipses in order to update a Kalman filter, which maintains the current estimate of the plane. The corresponding Kalman filter may be derived from the general form of

the equations given in [9] and are not given here due to limitations of space. The second grouping stage detects the axis of revolution by fitting a straight line through the centroids of the parallel planar ellipses, ensuring that ellipses whose centres are co-linear are also co-directional with the plane of the ellipses. (Since the ellipses we have detected are assumed to be from circular features, this implies that the cross-section of the solid of revolution is perpendicular to its axis.) For an axis of revolution, $R(\tau, \theta, D)$, the MD test is $R^t C_R^{-1} R < t$ (with three degrees of freedom) and C_R is the covariance of R .

The Mahalanobis distance/Kalman filter combination allows the grouping process to maintain the uncertainty associated with the ellipse estimation, and to propagate this uncertainty to the normal of the plane of the grouped ellipses and the axis of revolution.

The second part of our task is to match the detected (projected) circles with similar features in a 3D model. The inferred circles in the image provide position, viewpoint and scale information which can be directly used to map the model description via a $3D \Rightarrow 2D$ transformation onto the image feature data. Our intention is to find the rotation and translational parameters for the model-to-image transformation using a Kalman filter. The advantage of the iterative operation of the Kalman filter is that it allows the combination of evidence from different types of features (e.g. corners, circles, edges) in a consistent manner. Again, uncertainties associated with the measurement data are propagated through the transformation in order to generate projected model features which reflect the uncertainty in the measurement. As before, the fitting or matching process uses the Mahalanobis distance for selecting appropriate feature matches and a Kalman filter for refining the parameters of the viewpoint transformation.

4 Results

Figure 3 shows the detection and fitting of ellipses on a scene rich in projected circles. Edges are detected with a modified version of Shen's algorithm [10, 1] and are shown in 3b. Considerable noise in the detected image is derived in part from shadows, and from a high sensitivity setting for the detection algorithm. Following polygonal approximation and arc detection (detected arcs are shown in figure 3c), ellipses are constructed as detailed in section 3. Figure 3d shows the detected ellipses as estimated from the original line and arc data, and then augmented with locally connected and global line segments. Ellipses which are supported by only a small length of edge data in the image are ignored, since such ellipses are typically of high uncertainty, and will exhibit a poor fit. Ellipses have been grouped for several axes of revolution, though poorer detection performance on the toy cone has failed to group the resulting ellipses for that object. The axes and associated ellipses are shown overlaid onto the original image in figure 3e.

Figure 4 shows the same processing steps applied to a real scene (4a), taken from an industrial environment. As can be seen, the image is very complex, but still some useful ellipses are detected, and a grouping has been constructed. Figure 4b shows the edges detected, and 4c shows the ellipses. The final grouping is shown in figure 4d.

5 Conclusions and Further Work

It is well known that ellipse fitting is subject to a high uncertainty [5] and it is essential to take this into account for a 3D interpretation of such features in images. A combination of the Mahalanobis distance and Kalman filter, used throughout the interpretation process, allows this to be achieved in a consistent way.

The accuracy of the ellipse fitting is increased by grouping primitive image features (i.e. line segments) which may be part of the same ellipse. Multiple ellipses may then be grouped into parallel planar and co-axial sets, allowing the inference of 3D structures (e.g. solids of revolution) in the scene. As uncertainty has been propagated to these structures, this may be taken into account during subsequent matching processes with a model.

Although it is efficient to make use of arc detection to suggest possible ellipses in the image, this limits the detection of some ellipses, and has a tendency to miss important elliptical edge data. This is one aspect of the problem to which we are currently giving attention.

The problem of using the MD test is that when the uncertainty is high, in order not to select too many false features the threshold must be reduced, and it is not simple to find an optimal value. Current work is in hand to replace the MD test with a likelihood ratio (LR) test [1] in order to determine this optimal threshold.

In conclusion, dealing in a consistent way with the measurement error, and propagating uncertainty information through the low and intermediate processing stages provides robust support for matching against high-level symbolic object descriptions.

References

- [1] Brillault B., "A Probabilistic Approach to 3D Interpretation of Monocular Images", Ph.D. Thesis, City University (in preparation).
- [2] Ayache N., Faugeras O.D., "Building, registering and fusing noisy visual maps", *Proc. 1st ICCV*, pp. 73-82, London, England, 1987.
- [3] Lowe, D.G., "Perceptual Organisation and Visual Recognition", Kluwer, 1985.
- [4] Rosin P.L., West G.A.W., "Segmenting curves into elliptic arcs and straight lines", *Proc. 3rd ICCV*, Japan, 1990.
- [5] Porrill J., "Fitting ellipses and predicting confidence envelopes using a bias corrected Kalman filter", *Image and Vision Computing*, 8, pp. 37-41, 1990.
- [6] Rosin P.L., West G.A.W., "Perceptual grouping of circular arcs under projection", *Proc. BMVC90*, Oxford, UK, pp 379-382, 1990.
- [7] Grimson W.E.L., Huttenlocher D.P., "On the verification of hypothesized matches in model-based recognition" *Proc. 1st ECCV90*, pp 489-498, Antibes, France, 1990.

- [8] Masciangelo S., "3-D cues from a single view: detection of elliptical arcs and model-based perspective backprojection", **Proc. BMVC90**, pp. 223-228, Oxford, UK, 1990.
- [9] Kalman R. E., "A new approach to linear filtering and prediction problems", **Trans. ASME, J. of Basic Engineering** 35-45, March 1960.
- [10] Shen J., Castan S., "An optimal linear operator for edge detection", **Proc. CVPR'86**, June, 1986.




Figure 3: Grey image, edges, arcs, detected ellipses and grouped ellipses.




Figure 4: Grey level data, edges and detected and grouped ellipses in a real scene.
