Projectively Invariant Representations Using Implicit Algebraic Curves

David Forsyth, Joseph L. Mundy, Andrew Zisserman, and Christopher M. Brown[1]
Robotics Research Group
Department of Engineering Science
Oxford University
England

We demonstrate that it is possible to compute polynomial representations of image curves which are unaffected by the projective frame in which the representation is computed. This means that:

> The curve chosen to represent a projected set of points is the projection of the curve chosen to represent the original set.

We achieve this by using algebraic invariants of the polynomial in the fitting process. We demonstrate that our procedure works for plane conic curves. We show that for higher order plane curves, or for aggregates of plane conics, algebraic invariants can yield powerful representations of shape that are unaffected by projection, and hence make good cues for model based vision. Tests on synthetic and real data have yielded excellent results.

1 Introduction

It is common in machine vision to wish to represent a set of image points, $\mathcal{S} = (x_i, y_i)$, $i = 1 \cdots M$, by a polynomial curve, thus compressing its information content. This representation makes it possible to define a set of projectively invariant shape descriptors.

Representation and approximation are different goals. A curve that is a poor approximation of a set of points may serve as an excellent representation, if the procedure for choosing the curve is stable and is independent of the frame in which the curve is chosen. For machine vision, a good representation for a set of points has the frame independence property:

> Given an observation of the set in a transformed frame, the representation computed for this set is exactly the original representation transformed according to the change of frame.

A representation with this independence property need not be a good *approximation* to the data. This property is essential to guarantee that descriptive features of the curve are unaffected by image transformation.

Stability is a second important property of a representation. A stable representation has the property:

> A small change in the data will result in a small change in the representation.

[1]DAF acknowledges the support of Magdalen College, Oxford. AZ acknowledges the support of the Science and Engineering Research Council. JLM acknowledges the support of the General Electric Coolidge Fellowship. CMB acknowledges the support of the DARPA U.S. Army Engineering Topographic Laboratories Grant DACA76-85-C-0001 and the Air Force Systems Command (RADC, Griffiss AFB, NY) and Air Force OSR Contract F30602-85-C-0008, which supports the Northeast Artificial Intelligence Consortium.

This definition is meaningful for point sets and polynomial curves, because it is possible to specify what a small change is.

In this paper, we show that algebraic invariants can be used to achieve projectively invariant curve fitting, and that invariants can be useful in matching and identification tasks. This work considers the case where the set, $\mathcal{S}$, can be represented by a planar curve, C of known order. Furthermore, we assume that the "segmentation problem" has been solved, so that it is known which sets of points require distinct representations. The curve C is represented implicitly as a polynomial. This representation has the advantage that the invariant theory for such curves is well established, and that a projectively invariant error metric exists. This measure of error is called algebraic distance, and is described below.

The discussion in this paper is strongly focused on conics, because they are familiar curves which have been widely addressed in the vision literature. The mathematical development holds, however, for implicit polynomial curves and surfaces of any degree.

1.1 Background

Space does not permit a full review of the literature in this area. The interested reader is referred to [4]. A number of papers have contributed significantly to the line of thought presented here [1,2,11,12,13].

We write an implicit polynomial curve as $Q(x, y, \mathbf{p}) = 0$, where $\mathbf{p}$ is the vector of coefficients of the polynomial. In particular, for a conic we have $\mathbf{p} = [A\ B\ C\ D\ E\ F]$ and $Q(x, y, \mathbf{p}) = Ax^2 + Bxy + Cy^2 + Dx + Ey + F$. The *algebraic distance* of a point (x_i, y_i) from an implicit polynomial curve $Q(x, y, \mathbf{p})$ is $Q^2(x_i, y_i, \mathbf{p})$. The algebraic distance of a set of points from a curve is the mean of their individual algebraic distances. Algebraic distance is often used as a measure of the deviation of a point (x_i, y_i) from a polynomial curve because this error metric can lead to a linear system of equations for the curve fitting problem, where euclidean distance leads to complex, non-linear equations. Furthermore, when data can be viewed in different projections, euclidean distance is no longer a useful error measure, because a pair of points that are widely separated in one projection may be arbitrarily close together in a second. However, algebraic distance is not unique, since $kQ(x_i, y_i, \mathbf{p}) = 0$ defines the same polynomial curve. In order to make the value of $Q^2(x_i, y_i, \mathbf{p})$ unique, it is necessary to define a normalization function, $N(\mathbf{p})$ which is held constant during the fitting process. This constraint can be introduced by the use of Lagrange multipliers [2]. Note that if $N(\mathbf{p})$ is not a homogenous, positive (or negative) definite function (i.e. one that is always positive, and zero only if the argument is zero) of $\mathbf{p}$, then it may not be possible to fit certain curves. In particular, if we fit a curve requiring $N(\mathbf{p}) = 1$, if $N(\mathbf{p})$ is not positive (or negative) definite, then there will be some $\mathbf{p}$ such that $N(\mathbf{p}) = 0$, and this curve is clearly unattainable by our fitting process. This effect can cause poor fits.

In this paper, we show that normalizing by an algebraic invariant (defined below) of the group of changes of frame, means that the fitted curve has the frame independence property described above. Furthermore, we demonstrate invariant curve fitting leading to curve descriptors that are unaffected by the frame in which the curve is viewed, and demonstrate frame invariant curve descriptors.

2 Algebraic invariants

An invariant of a transformation is defined as follows:

> **Definition** An invariant, $I(\mathbf{p})$, of a function $f(\mathbf{x},\mathbf{p})$ subject to a group, $\mathcal{G}$, of transformations acting on the coordinates $\mathbf{x}$, is transformed according to $I(\mathbf{p}') = I(\mathbf{p})h(g)$. Here $g \in \mathcal{G}$ and $h(g)$ is a function only of the parameters of the transformation and does not depend on the coordinates, $\mathbf{x}$, or on the parameters, $\mathbf{p}$. $I(\mathbf{p})$ is a function only of the parameters, $\mathbf{p}$.

For linear coordinate transformations, $\mathbf{x} = \mathbf{T}\mathbf{x}'$, the form of the invariance relation becomes $I(\mathbf{p}') = I(\mathbf{p})|\mathbf{T}|^w$ where $\mathbf{T}$ is the transformation matrix and $|\mathbf{T}|$ indicates the determinant of $\mathbf{T}$. In this case, it can be shown that $\mathbf{p}' = \mathcal{T}_g\mathbf{p}$, where $\mathcal{T}_g$ is a transformation matrix for the polynomial coefficients corresponding to the group element g. The exponent, w, is called the weight of the invariant. An invariant with $w = 0$ is called a *scalar* invariant. The function $h(g)$ is represented by the determinant of the transformation matrix. The simplest example of an algebraic invariant is the discriminant of a quadratic equation of one variable, which is an invariant of weight 2, under translation and scaling. Algebraic invariants formed a major research topic of 19th century mathematics: an introduction can be found in, for example, [6]. A modern perspective of some of this work is given by [8], or by [14].

A conic polynomial can be expressed as:

$$Q(\mathbf{x},\mathbf{p}) = \mathbf{x}^T\mathbf{P}\mathbf{x}, \text{ where } \mathbf{P} = \begin{bmatrix} A & \frac{B}{2} & \frac{D}{2} \\ \frac{B}{2} & C & \frac{E}{2} \\ \frac{D}{2} & \frac{E}{2} & F \end{bmatrix} \tag{1}$$

$\mathbf{P}$ is the coefficient matrix, and $\mathbf{x} = [x_1, x_2, x_3]^T$. Note that the standard expression for the conic in euclidean coordinates is obtained by performing the indicated matrix operations and then setting $x_3 = 1$. In a different frame, where $\mathbf{x} = \mathbf{T}\mathbf{x}'$, this conic will be represented by a different matrix, $\mathbf{P}'$. However, the value of the conic at some point on the plane is fixed, whatever the coordinates of this point, so $Q(\mathbf{x},\mathbf{p}) = \mathbf{x}^T\mathbf{P}\mathbf{x} = Q(\mathbf{x}',\mathbf{p}') = \mathbf{x}'^T\mathbf{P}'\mathbf{x}' = \mathbf{x}'^T\mathbf{T}^T\mathbf{P}\mathbf{T}\mathbf{x}'$, and this is true for all $\mathbf{x}$, so that $\mathbf{P}' = \mathbf{T}^T\mathbf{P}\mathbf{T}$.

In effect, $\mathbf{T}^T$ and $\mathbf{T}$ "strip" the effect of the change in coordinate system, so that for a given point on the projective plane, although the point has different coordinates in the different coordinate system, the value of the polynomial is the same. For transformations corresponding to general matrices in homogenous coordinates (i.e. projective transformations combined with a choice of scale), it is easily seen that $|\mathbf{P}|$ is an invariant of weight 2. This is in fact the only invariant of a single conic under projective transformations.

3 Normalization by invariants

Bookstein [2] considers using algebraic distance as an error measure, and shows that, in this case, if one normalises a curve by a *scalar* invariant of the transformation group that will act on it, the curve chosen to fit a set of points will be unaffected by the frame in which the curve is chosen. In fact a stronger theorem is possible: it is sufficient to use an invariant *of any weight.*

Theorem 1 Let $I(\mathbf{p})$ be an invariant of the polynomial form $Q(\mathbf{x}, \mathbf{p})$ under a group of linear transformations $\mathcal{G}$. Assume I is homogenous of degree n, with weight w. Let $<\mathbf{p}>$ be the parameter vector determined by minimizing $\sum_i Q^2(\mathbf{x}_i, \mathbf{p})$ over a set of points, $\mathbf{x}_i$, subject to the constraint, $N(\mathbf{p}) = I(\mathbf{p}) = constant$. If the point set is transformed under $\mathcal{G}$, i.e., $\mathbf{x} = \mathbf{T}_g\mathbf{x}'$, let $\mathcal{T}_g$ be the corresponding transformation matrix *for the coefficients* $\mathbf{p}$. The coefficients of the polynomial fitted to the point set in the new frame are given by $<\mathbf{p}'>$. Assume that n is odd or that w is even (or both). Under these conditions, we have:

$$<\mathbf{p}> = k_g.\mathcal{T}_g <\mathbf{p}'>$$

where k_g is a scalar depending on $g \in \mathcal{G}$.

The proof of this theorem is given in appendix A. The curve chosen by the approximation process is then effectively decoupled from the frame in which we make observations, and has the desired frame independence property. Note that this property is true of the *curve*, $\{\mathbf{x} \mid Q(\mathbf{x}, \mathbf{p}) = 0\}$, and not of the *polynomial* $Q(\mathbf{x}, \mathbf{p})$. It is not difficult to convince oneself that this theorem must be true for the case where the normalization is a scalar invariant, where the algebraic distance of a curve from a set of observations is unaffected by an admissible transformation. Intuitively, it is true for the case of an invariant of any weight because, although the algebraic distance varies for an admissible transformation, the *curve* with minimum algebraic distance remains the same.

4 Fitting conics using algebraic distance

The algebraic distance from a set of points to a curve is:

$$\overline{Q^2(\mathbf{x}_i, \mathbf{p})} = \frac{1}{M}\sum_{i=1}^{i=M}\left(\mathbf{x}_i^T\mathbf{P}\mathbf{x}_i\right)^2 = \mathbf{p}^T\mathbf{S}\mathbf{p}$$

with $\mathbf{P}$ the conic coefficient matrix, $\mathbf{p}$ the vector of conic coefficients, M the number of data points. The sum is over all the data points. $\mathbf{S}$ is referred to as the scatter matrix for the conic, and its form can be explicitly reconstructed from the above expression. Note that the scatter matrix is positive, and unless an exact fit is possible, positive definite.

The error is minimized subject to the normalizing constraint, $N(\mathbf{p}) = constant$, where $N(\mathbf{p})$ is the determinant of the scatter matrix. It can be shown that for this particular problem, the value of this constant is not important [4]. We chose to use $N(\mathbf{p}) = 1$ throughout this work.

It is possible to show that a global minimum exists by demonstrating that this problem is equivalent to maximising a cubic form confined to the five dimensional sphere, which is compact [4]. Extensive experiments with conventional iterative methods, which halt at local extrema, suggest that the function may not always be convex. These methods perform poorly as a result, because this technique requires a global extremum.

We use a method that evaluates the function on increasingly fine subdivisions of the five-sphere, and retains the largest value found. This method avoids combinatorial explosion by pruning the search space, using the fact that there is an upper bound, say B, on the norm of the gradient of the objective function. In particular, if the most recent estimate of the global maximum exceeds the value at the centre of a cell whose radius is r by Br, then the cell need not be subdivided further because the cell cannot contain the global maximum. At each subdivision, because the domain of the function under consideration has been pruned, the upper bound on the size of the gradient is recomputed. This process leads to an algorithm that in practice finds very good global extrema reasonably fast, and in practice appears to avoid the worst effects of combinatorial explosion. A detailed discussion of this method appears in [5].

5 Results

The invariant fitting procedure can be checked for frame independence by comparing fits achieved in different frames. This procedure involves constructing a data set, fitting to that set in one frame, projecting the set, fitting in a second frame, and comparing the fit in the second frame with the projection of the fit in the first frame. Any discrepancy is due to a dependence of the fitting procedure on the frame in which the fit is achieved.

Data sets lying on conics and scattered data sets were tested in this way. We do not show the results for data sets drawn from a conic, because any fitting procedure based on algebraic distance that does not have an unreasonable normalisation will be frame independent on such data sets. A result for a scattered data set is shown in figure 1. In particular, the fitting process was wholly unaffected by the frame in which the curve was fitted.

Stability can be investigated by adding gaussian noise to the data sets *in each frame.* Thus, in the example shown, the data set is projected into the appropriate frame, and noise is added after projection. This means that the data sets are no longer exactly within a projection of one another. One then fits a curve to the noisy data in each frame, and checks that the curves are within projection of one another. Using this procedure gives a good indication of performance in the presence of quantization and imaging noise. From figure 2, our technique appears to degrade gracefully in the presence of limited amounts of gaussian noise. These tests suggest that the fitting technique is frame independent and reasonably stable. A measure of the stability can be obtained by computing the joint scalar invariants of a pair of conics, which are described below. These joint invariants are nearly constant for pairs of conics fitted to different views of the same pair of curves (see table 1). These results come from image data, where noise effects mean that the data sets are no longer within projection of one another.

5.1 Joint scalar invariants of conics

A pair of conics admits two joint scalar invariants, described in detail in [4] or in [16]. For two conics, Q_a and Q_b, in addition to $I_{a3} = |\mathbf{P}_a|$ and $I_{b3} = |\mathbf{P}_b|$, there are the joint invariants of weight 2,

$$I_{ab1} = \sum_i \sum_j {P_a}^{ij} P_{bij}, \quad I_{ab2} = \sum_i \sum_j {P_b}^{ij} P_{aij}$$

where $\mathbf{P}_a$ and $\mathbf{P}_b$ are the coeficient matrices of Q_a and Q_b respectively. ${P_a}^{ij}$ and ${P_b}^{ij}$ are the cofactors of the corresponding matrix elements. Two independent scalar invariants for a configuration of two conics can then be formed by,

$$I_{ab3} = \frac{I_{ab1}}{|\mathbf{P}_a|}, \quad I_{ab4} = \frac{I_{ab2}}{|\mathbf{P}_b|}$$

Conics	First joint invariant	Second joint invariant
Conics a and b from figure 3a	3.419	3.546
Conics a and b from figure 3b	3.418	3.543
Conics a and b from figure 3c	3.414	3.538
Conics a and b from figure 3d	3.407	3.528

Table 1: The joint scalar invariants computed for conics a and b for the four different images of the tape from different positions and angles, shown in figure 3. Note that the joint invariants for the coplanar conics a, b from the four images are effectively constant.

Conics	First joint invariant	Second joint invariant
Conics a and c from figure 3c	-7.547	-6.359
Conics a and c from figure 3d	-9.144	-6.778

Table 2: The joint scalar invariants computed for the non-coplanar conics a and c in figure 3c and figure 3d. Because these curves are not coplanar, the joint scalar invariants change as the viewpoint changes.

For two sets of coplanar data points, the joint scalar invariants should be constant whatever the camera viewpoint. Figure 3 shows fitted curves superimposed on image data. Table 1 shows the joint invariants for these curves, observed in different images from different angles. The joint invariants for this pair of coplanar curves, observed from different viewpoints in different images, agree very closely. However, if one data set is not coplanar with a second, then the joint scalar invariants computed should vary with the camera position, because they are invariant only under planar projection. This yields an elegant test for coplanarity for sets of points, demonstrated in table 2.

Figure 4 shows an instance of a model found in a complex scene by fitting conics to all the curves in that scene, and then marking those curves in the scene that have the correct pair of joint scalar invariants. These results argue that the invariant fitting technique is successful and stable in computing projectively invariant representations for pairs of curves. The test data shown in the figures, and other test sets, can be obtained from the first author.

6 Discussion and conclusions

This work has raised a number of issues that remain open. Our work leads us to believe that stability of fit and frame independence are the most important properties of a representation to be used for recognition in vision applications. Although the invariant representation we used is frame independent, we do not yet understand the manner in which uncertainties resulting from, for example, quantization error, reflect in the invariants. Furthermore, because the fitting process chooses a curve that best represents the entire dataset, it is sensitive to occlusion. It is not clear at this stage how to build a projectively invariant representation of a curve that is insensitive to occlusion.

Our results indicate that the combination of algebraic distance and invariant normalization leads to representations of data sets that are unaffected by projection. Theorem 1 shows that this will be true for curves of a higher order than conics. Projectively invariant representations have major applications. Three examples appear below.

In recovering three dimensional object contours by integrating multiple image views, it is important that curve representations are transformationally consistent so that object boundaries seen from different viewpoints can be matched.

We have demonstrated that algebraic invariants may be used to describe curves by projectively invariant signatures (e.g. the joint scalar invariants for pairs of conic curves). Weiss [16] first proposed using these invariants for vision, but computing such invariant signatures successfully requires an invariant fitting procedure for conics such as the one we have developed. The invariance of the ratio of areas of concentric squares and concentric triangles, used by Nielson [10] for robot navigational landmarks follows from the conic pair invariants by considering concentric circles. Our results do not require the invariants to be defined using the point cross ratio, which Nielson used in his analysis.

One of the most important issues in model-based object recognition is efficient indexing of image features onto corresponding three dimensional model features. The bulk of current research is focussed on polyhedral models and image features comprised of groups of image edges and vertices, [7,9,15]. As figure 4 indicates, algebraic invariants present a potentially efficient and reliable index between two dimensional image curves and corresponding models.

Although we have demonstrated our techniques with conics, higher order plane curves have richer sets of invariants. For example, a single cubic has two invariants, from which a single scalar invariant can be computed. Thus, a projectively invariant shape description can be computed for a single cubic curve, while for conics two curves are necessary. The techniques we have shown here are relevant in their present form only for coplanar curves. Geometrical considerations suggest that, for example, rigidly coupled pairs of plane curves, subjected to euclidean actions and then projected, admit scalar invariants, meaning that there is ample scope for these ideas to be extended.

Acknowledgements

Our thanks to Michael Brady and the Robotics Research Group, Oxford, for providing a pleasant environment in which to think and work. Our thanks to Andrew Blake for a number of stimulating discussions. Thanks to GE for the phone calls that made this paper possible.

References

[1] Agin, G.J. "Fitting ellipses and general second order curves," Carnegie Mellon University CMU-RI-TR-81-5, undated.

[2] Bookstein, F. "Fitting conic sections to scattered data," *CVGIP*, **9**, 56-91, 1979.

[3] Canny, J.F. "Finding edges and lines in images," TR 720, MIT AI Lab, 1983.

[4] Forsyth, D.A., Mundy, J.L., Zisserman, A.P. and Brown, C.M. "Projectively invariant data representation by implicit algebraic curves: a detailed discussion," Internal report 1814/90, Department of Engineering Science, Oxford University, 1990.

[5] Forsyth, D.A. "Finding the global extremum of a function with bounded derivative over a compact domain," Internal report 1815/90, Department of Engineering Science, Oxford University, 1990.

[6] Grace, J.H. and Young, A. *The algebra of invariants*, Cambridge University Press, Cambridge, 1903.

[7] Grimson, W.E.L. and T. Lozano-Pérez, "Localizing overlapping parts by searching the interpretation tree," IEEE T-PAMI, **PAMI-9**, 4, 469-481, 1987.

[8] Kung, J.P.S. and Rota, G.-C. "The invariant theory of binary forms," *Bull. Amer. Math Soc.*, **10**, 27-85, 1984

[9] Lowe, D. *Perceptual Organisation and Visual Recognition*, Kluwer, Boston, 1985.

[10] Nielson, L. "Automated guidance of vehicles using vision and projectively invariant marking," *Automatica*, **24**, 2, 135-148.

[11] Porrill, J. "Fitting ellipses and predicting confidence envelopes using a bias corrected Kalman filter," Proc. 5'th Alvey vision conference, 1989.

[12] Pratt, V. "Direct least-squares fitting of algebraic surfaces," *ACM SIGGRAPH*, **21**, 145-151, 1987.

[13] Sampson, P.W. "Fitting conic sections to "very scattered" data: an iterative refinement of the Bookstein algorithm," *CVGIP*, **18**, 97-108, 1982.

[14] Springer, T.A. *Invariant theory*, Springer-Verlag lecture notes in Mathematics, **585**, 1977.

[15] Thompson, D.W. and J.L. Mundy, "3D model matching from an unconstrained viewpoint," *Proc. IEEE Conf. on Robotics and Automation*, 1987.

[16] Weiss, I. "Projective invariants of shapes," *Proc. DARPA IU workshop*, 1125-1134, 1989.

Appendix: The invariant fitting theorem

Theorem 1 Let $I(\mathbf{p})$ be an invariant of the polynomial form $Q(\mathbf{x},\mathbf{p})$ under a group of linear transformations $\mathcal{G}$. Assume I is homogenous of degree n, with weight w. Let $<\mathbf{p}>$ be the parameter vector determined by minimizing $\sum_i Q^2(\mathbf{x}_i,\mathbf{p})$ over a set of points, $\mathbf{x}_i$, subject to the constraint, $N(\mathbf{p}) = I(\mathbf{p}) = constant$. If the point set is transformed under $\mathcal{G}$, i.e., $\mathbf{x} = \mathbf{T}_g\mathbf{x}'$, let $\mathcal{T}_g$ be the corresponding transformation matrix *for the coefficients* $\mathbf{p}$. The coefficients of the polynomial fitted to the point set in the new frame are given by $<\mathbf{p}'>$. Assume that n is odd or that w is even (or both). Under these conditions, we have:

$$<\mathbf{p}>= k_g.\mathcal{T}_g <\mathbf{p}'>$$

where k_g is a scalar depending on $g \in \mathcal{G}$.

Proof: With $\mathbf{p}^T = [A\ B\ C\ D\ E\ F]$, we have $\overline{Q^2(\mathbf{x}_i,\mathbf{p})} = \mathbf{p}^T\mathbf{S}\mathbf{p}$, where S is the positive definite scatter matrix defined above. $<\mathbf{p}>$ is determined by minimizing the mean algebraic distance subject to the constraint of constant normalization. The number of polynomial coefficients depends on the degree r of Q, and will be denoted by $M(r)$. For example, for a conic $M(2) = 6$ and for a cubic, $M(3) = 10$. The appropriate Lagrangian is $\mathcal{L} = \mathbf{p}^T\mathbf{S}\mathbf{p} + \lambda N(\mathbf{p})$

At any local minimum of mean squared algebraic distance $\nabla_{\mathbf{p}}\mathcal{L} = \mathbf{0}$. This gives $M(r)$ equations in $M(r) + 1$ unknowns. The final equation is given by $N(\mathbf{p}) = constant$. Suppose that the point set is transformed by some element, $g \in \mathcal{G}$ so that $\mathbf{x} = \mathbf{T}_g\mathbf{x}'$. The polynomial coefficients are transformed linearly by a $M(r)\times\ M(r)$ transformation matrix, $\mathcal{T}_g$, such that $\mathbf{p}' = \mathcal{T}_g\mathbf{p}$, and the scatter matrix transforms to $\mathbf{S}'$.

The form of $\mathbf{S}'$ can be determined by observing that for any polynomial with coefficients $\mathbf{p}$ in some frame, the algebraic distance at a given point on the plane is not affected by the particular frame in which the point is expressed. Thus

$$\mathbf{p'}^T\mathbf{S}'\mathbf{p}' = \mathbf{p}^T\mathcal{T}_g^T\mathbf{S}'\mathcal{T}_g\mathbf{p} = \mathbf{p}^T\mathbf{S}\mathbf{p}$$

This is true for all $\mathbf{p}$, so that $\mathbf{S} = \mathcal{T}_g^T\mathbf{S}'\mathcal{T}_g$.

The normalization is an invariant under $\mathcal{G}$ so $N(\mathbf{p}') = |\mathcal{T}_g|^w N(\mathbf{p})$, where w is the weight of the invariant. Applying these transformation rules to the Lagrangian,

$$\mathcal{L} \;=\; \mathbf{p}^T\mathbf{S}\mathbf{p} + \lambda N(\mathbf{p}) \;=\; \mathbf{p}^T\mathcal{T}_g^T\mathbf{S}'\mathcal{T}_g\mathbf{p} + \lambda\frac{1}{|\mathcal{T}(g)|}^{w} N(\mathcal{T}_g\mathbf{p}) \;=\; \mathbf{p}'\mathbf{S}\mathbf{p}' + \lambda' N(\mathbf{p}') \;=\; \mathcal{L}'$$

Hence the Lagrangian for $\mathbf{p}'$ is the same as the Lagrangian for $\mathcal{T}_g\mathbf{p}$. It follows that $<\mathbf{p}'>$ can differ from $\mathcal{T}_g<\mathbf{p}>$ by at most a scalar factor, $k(g)$. The final equation, $N(<\mathbf{p}>) = constant$, can be satisfied by scaling $<\mathbf{p}>$. This will fail only when N is homogenous of even degree, and the signs of $N(<\mathbf{p}>)$ and $N(\mathcal{T}_g<\mathbf{p}>)$ differ. Thus, the theorem is proven.

List of Figures

1. The diagonal crosses show a scattered data set. The crosses show that set in a different projection. Curve *a* was fitted to the diagonal crosses. Curve *b* was fitted to the crosses, and curve *c* is the appropriate projection of curve *a*. The two curves are indistinguishable at the resolution of laser-printer output. This demonstrates the frame independence of the fitting process.

2. The crosses show the projected data set of figure 1, and line up with the crosses in figure 1. Curve *a* was fitted to the crosses. Isotropic gaussian noise of $\sigma = 3$ was added to this set *in this frame*, to give the points plotted as diagonal crosses. The data set is thus no longer a projection of the original set. Curve *a* is the curve fitted to the diagonal crosses, curve *b* is the projection of curve *a* of figure 1. These curves now differ slightly, as a result of the noise. The small difference indicates a degree of stability for fitting circles and ellipses, and suggests that the frame independence property is stable.

3. Four images of a computer tape, with two fitted conics in overlay. The data for the conics in these images was obtained by acquiring the image edges using a local implementation of Canny's [3] edge finder, linking edges, and then choosing corresponding curves by hand. In these images, the conics have been drawn three pixels thick to make them visible. These conics were used to obtain the joint scalar invariants shown in tables 1 and 2.

4. The joint scalar invariants of a pair of conics can be used to find instances of models in scenes, when the objects involved have plane conic curves which lie on their surfaces. Here we show an instance of a computer tape found in a cluttered scene by fitting conics to all of the curves, and marking those pairs of conics with the correct joint scalar invariants. The data for the conics in this image was obtained by acquiring the image edges using a local implementation of Canny's [3] edge finder and then linking these edges.

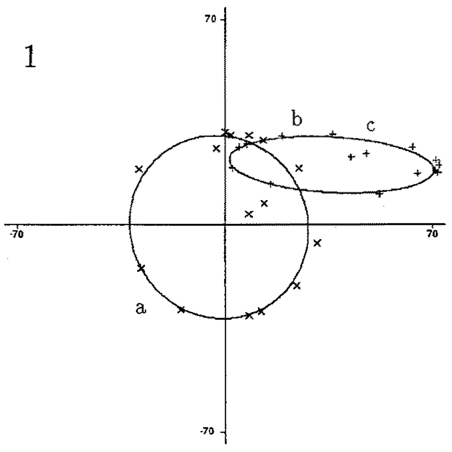
1
70
-70
70
-70
a
b
c

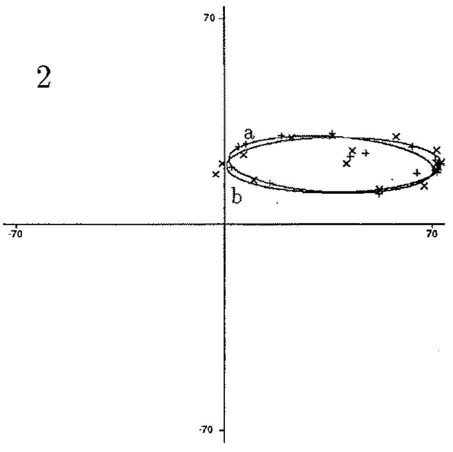
2
70
-70
70
-70
a
b

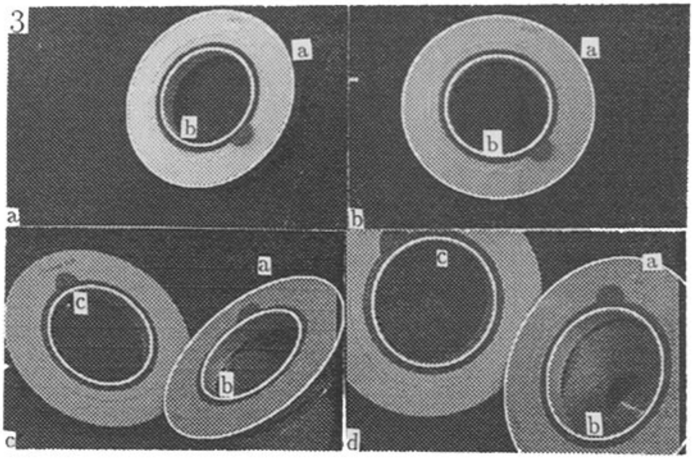
3
a
b
a
b
a
b
c
a
b
c
a
b
d

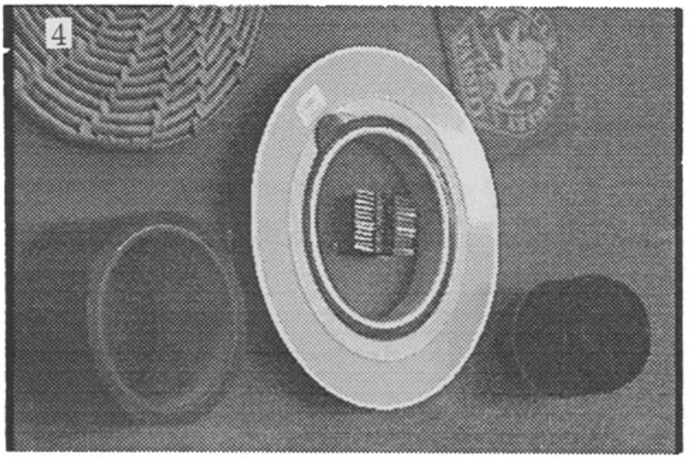
4