PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Shape from Symmetry: Detecting and Exploiting Symmetry in Affine Images

Dipti Prasad Mukherjee, Andrew Zisserman and Michael Brady

Phil. Trans. R. Soc. Lond. A 1995 **351**, 77-106
doi: 10.1098/rsta.1995.0026

Shape from symmetry: detecting and exploiting symmetry in affine images

By Dipti Prasad Mukherjee, Andrew Zisserman and Michael Brady

Robotics Research Group, Department of Engineering Science, University of Oxford, Oxford OX1 3PJ, U.K.

Contents

We investigate the constraints placed on the image projection of a planar object having local reflectional symmetry. Under the affine approximation to projection, we demonstrate an efficient (low-complexity) algorithm for detecting and verifying symmetries despite the distorting effects of image skewing. The symmetries are utilized for three distinct tasks: first, determining image back-projection up to a similarity transformation ambiguity; second, determining the object plane orientation (slant and tilt); and third, as a test for non-coplanarity amongst a collection of objects. These results are illustrated throughout with examples from images of real scenes.

Phil. Trans. R. Soc. Lond. A (1995) **351**, 77–106
Printed in Great Britain

1. Introduction

The representation and recognition of shapes by computer has numerous applications: as a step towards automating processes such as inspection; acquisition of objects from a conveyor belt or container; reconnaissance; and, navigation by an autonomous system. Also, such work can potentially contribute to understanding one of the most exquisite and effortless of human competences: the rapid recognition of familiar shapes even when they are partially occluded by others, when their surface colour or texture is unfamiliar, and when they are viewed from a wide range of vantage points. From a very early age, we can learn new classes of shapes, learn to discriminate subclasses, and then mobilize those new representations to effect recognition.

Understanding how to represent and recognize shapes has, however, proved to be a remarkably difficult task, both for computer vision and for perceptual psychology. So much so, that the current state of the art is that only limited classes of shapes can be recognized reliably from a limited range of poses. To date, perceptual psychology has been of limited usefulness, for though theories of shape abound (Beiderman 1987; Corballis 1988), they are too vaguely formulated to be implemented in a computer recognition system. On the other hand, most computer vision approaches to shape have either emphasized gross shape characteristics (for example, low-order moments or the first few coefficients of the Fourier or other transform of the contour function of the shape), or have relied on highly localized features, such as an estimate of points of high curvature along the bounding contour. Gross shape representations have insufficient discriminatory power and are sensitive to occlusion, while very local representations are subject to measurement noise that is unavoidable in practice.

The most advanced representation and recognition techniques developed to date in computer vision have explored representations intermediate between these extremes, and have exploited one or more of: relational constraints between parts of a shape, prior models, symmetry properties of the shape, or affine/projective invariance. The relational constraints approach relies upon precisely known algebraic relationships between different parts of a shape and has only been explored for the case of polyhedra, albeit classes of polyhedra defined parametrically (Grimson 1990; Reid 1991). Model-based recognition, even for non-polyhedral shapes, has enjoyed some success, but leaves aside the question of how recognition is effected when models are not available (Grimson 1990; Lowe 1985; Reid 1991). The present paper is a contribution to the symmetry and affine/projective invariance approaches to shape representation.

Many important classes of shapes, from faces and leaves through to manufactured items such as many stamped metal parts and profiles of aeroplanes and buildings exhibit one or more symmetries. The essential idea of a symmetry is a motion (Field & Golubitsky 1992): 'suppose you have an object and pick it up, move it around and set it down. If it is impossible to distinguish between the object in its original and final positions, we say that it has a symmetry'. Within the plane, symmetry comprises reflections, rotations and translations, in which latter case symmetry corresponds to pattern repetition. The line of thinking encapsulated by the quotation leads inexorably to modelling symmetry using the operations of group theory (Weyl 1952), a point to which we return when discussing affine/projective invariant representations of shape. However, requiring the transformed object to be 'impossible to distinguish' from the original is

far too restrictive both for computer vision and for human perception. Real objects such as faces, pears, wrenches and the outlines of fish (Strachan 1993) are only *approximately* symmetric and, more significantly, the shape only exhibits symmetries locally between segments of a shape or pattern.

Over the past thirty years, researchers in computer vision have explored a number of aspects of symmetry to generate symbolic representations of shape, culminating in the systems reported by Blum & Nagel (1978) and Connell & Brady (1987) and, more recently, by Rom & Medioni (1993). The first exploration of local support for symmetry was by Blum (1973) in his study of representations of biological shapes to effect recognition, determine abnormalities and monitor growth. The idea of local symmetry was put on a more solid mathematical footing by Giblin & Brassett (1985), who defined the *symmetry set* of a shape as the locus of the centres of all circles bitangent to a shape's bounding contour. A number of algorithms have been developed for computing the loci of local symmetries of shapes, particularly reflectional (Brady & Assada 1984; Scott *et al.* 1989) and rotational (Fleck 1985) symmetries, and they have been demonstrated to work reliably on a range of shapes, generating representations useful for recognition.

However, all such algorithms and representations share a severe limitation: symmetry is not preserved under skew, corresponding to the shape being viewed other than in a fronto-parallel plane. Simply stated, symmetry axes computed in an image of a shape taken from a non-fronto-parallel vantage point are not in general the transformed fronto-parallel symmetry axes. Despite this mathematical inconvenience, 'skewed symmetries' (that is, reflectional symmetries viewed from a non-fronto-parallel vantage point) such as those shown in figure 1 strongly suggest actual symmetries and constrain the plane in which they are perceived to lie. The mathematical fact that the skew symmetry may be an accident of projection is evidently discounted. Indeed, Wagemans (1993) has recently provided evidence that skewed symmetry is a non-accidental property of a shape that the human visual system exploits. Kanade (1981) was the first to analyse mathematically symmetries skewed by image projection, and proposed heuristics to interpret a skew symmetry as a real symmetry viewed from some (unknown) direction, which he represented using gradient space. Van Gool *et al.* (1992) proposed arc length space (ALS) to extract affine symmetry information using semi-differential invariants.

Notwithstanding the cogency of skewed symmetric shapes, and the information they seem to present about the pose (rotation and translation relative to the camera) of the shape, the fact remains that symmetry is not invariant under projection. This has recently persuaded a number of authors to abandon temporarily such representations to explore projective and affine invariants of two-dimensional shapes (Mundy & Zisserman 1992; Rothwell *et al.* 1992*a*). Like symmetry, invariance is also normally formalized using the mathematics of group theory, particularly Lie group theory. Systems have been developed that exploit a variety of invariants to recognize overlapped shapes viewed from a variety of poses. This substantial progress has been bought at the cost of regressing shape representations, now in a 'canonical frame', to descriptions invariant of pose.

In essence, the present paper is a step towards reconciling symmetry and invariance. For the present, we concentrate on local reflectional symmetries of smoothly curved planar objects, though the methods are equally applicable to polygonal objects, and in other work we have extended the ideas to certain three-dimensional

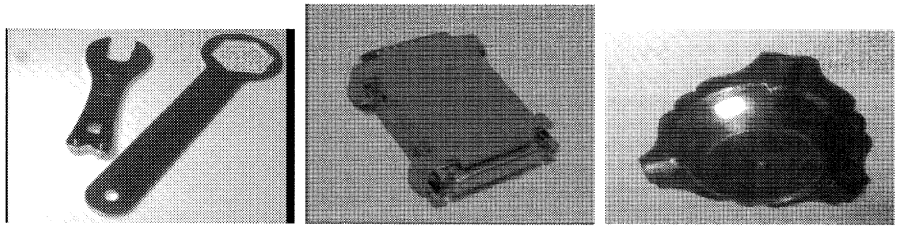

Figure 1. Examples of skewed symmetry. The objects differ in the number and disposition of the symmetry axes: (*a*) each object has a single bilateral symmetry; (*b*) the object has two orthogonal bilateral symmetries; and (*c*) the object has three bilateral symmetries, but the symmetry axes are not orthogonal. In each case the symmetry is global, and *on the object* lines between points related by the reflectional symmetry are orthogonal to the axis of reflection. However, in the image these lines and the axis are not orthogonal, in general, but skewed.

shapes (Fawcett *et al.* 1993). Suppose then that a planar object has a bilateral symmetry; how does this constrain its image projection? If the two 'sides' of the contour have a mirror symmetry, then one can be transformed onto the other by a reflection. A reflection is a particular type of affine transformation. We assume that the projection between the object and image planes can also be approximated by an affine transformation. It is easy to show, for example using a Taylor-series expansion of the projection equations of a pinhole camera, that this is a very good approximation provided the field of view is small, and the range of depths encompassing the object is at least an order of magnitude less than the distance of the object from the camera. This is satisfied in many practical cases. Extensions to (an exact) projective transformation for planar objects, and to symmetries of three-dimensional objects are discussed in § 5.

Consequently, since affine transformations form a group, the transformation between the two sides of the contour *in the image* is affine. This immediately provides an algorithm, albeit one that is computationally expensive, for detecting possible symmetries: if two image contours can be mapped onto each other by an affine transformation (six degrees of freedom), then the object could have had a reflectional symmetry. A key result of the present paper is given in § 2, where it is shown that the image transformation is actually a subset of the affine transformations with only *three* degrees of freedom. This, and the use of affine index functions, is used to develop an efficient algorithm for detecting and verifying symmetries (§ 3).

The mathematical framework is established in § 2. It is shown that a single symmetry is sufficient to unskew the back-projection, to give a one-parameter family of symmetric shapes that could have given rise to the image. A second coplanar symmetry is sufficient, in general, to uniquely determine the aspect ratio of the plane. This determines the back projection up to a similarity transformation (translation, rotation and isotropic scaling). This is achieved without any knowledge of the intrinsic camera parameters, an important consideration in practice since the automatic calibration of camera intrinsic parameters is a poorly conditioned nonlinear problem (Tsai 1986). If, however, the camera aspect ratio is known, the back projection determines the slant (up to a two-fold ambiguity, often referred to as the Necker ambiguity) and tilt of the object plane. Finally, a test for non-coplanarity is given for two symmetric objects from a single image.

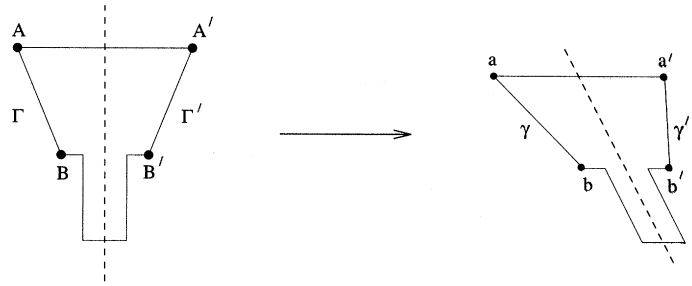

Figure 2. γ and γ' are the affine images of two corresponding sides of a planar object, Γ and Γ' respectively, with bilateral symmetry. Affine transformations preserve parallelism, so lines joining corresponding points in the image, for example the lines aa′ and bb′, are parallel. Affine transformations also preserve length ratios on parallel lines. In particular midpoints are preserved, so the imaged symmetry axis passes through the midpoints of aa′ and bb′.

Throughout the paper, results are shown from an implementation of the theory outlined in §§ 2 and 3. The reader may care to pause at this point to view figures 14, 15; figures 16, 18; and figures 19, 20, which show typical results using the implemented program.

2. Mathematical framework

In this section we study the constraints on the transformation between two image contours if they are the projections of corresponding sides of a planar object with bilateral symmetry. Figure 2 illustrates the situation under consideration: the image contours γ and γ' are the images, assumed to be affine, of two corresponding sides, Γ and Γ' respectively, of a planar object with a bilateral (mirror) symmetry. A key attribute of affine transformations is that they preserve parallelism, so lines joining corresponding points in the image, for example the lines aa′ and bb′, are parallel. Affine transformations also preserve length ratios on parallel lines. In particular, midpoints are preserved, so the imaged symmetry axis passes through the midpoints of aa′ and bb′.

(*a*) *Image transformation*

The following theorem explores the image transformation between the affine images γ and γ' in more detail. In particular, it fixes notation that will be used subsequently to unskew images.

Theorem 2.1. *Suppose two curves γ and γ', as in figure 2, are the images of two corresponding sides of a planar object with bilateral symmetry. Suppose further that image projection can be represented by an affine transformation.*

Then the transformation between γ and γ' has the following properties.

Property 1. γ and γ' are related by an affine transformation. That is, if $\boldsymbol{x}$ is a point on γ then there is a point $\boldsymbol{x}'$ on γ' such that

$$\boldsymbol{x}' = A\boldsymbol{x} + \boldsymbol{b}, \tag{2.1}$$

where A is a non-singular 2×2 matrix, and $\boldsymbol{b}$ is a two vector.

Property 2. The affine transformation $\{A, \boldsymbol{b}\}$ satisfies the following constraints:

(i) $A^2 = I$;

(ii) $|A + I| = 0$; *and*
(iii) $|A - I| = 0$.

Property 3. The matrix A *has eigenvectors* $\boldsymbol{a}$ *and* $\boldsymbol{b}$ *(*$\boldsymbol{b}$ *as above) with eigenvalues* $+1$ *and* -1, *respectively. Vector* $\boldsymbol{a}$ *is parallel to the symmetry axis, vector* $\boldsymbol{b}$ *is parallel to* $\boldsymbol{x}' - \boldsymbol{x}$.

Property 4. The transformation has three degrees of freedom. It can be parametrized by a, b_x, b_y, *where*

$$A = \begin{bmatrix} a & -b_x(1+a)/b_y \\ -b_y(1-a)/b_x & -a \end{bmatrix}, \quad \boldsymbol{b} = \begin{bmatrix} b_x \\ b_y \end{bmatrix}. \tag{2.2}$$

Property 5. With this notation, the image (skewed) symmetry line is

$$(1-a)b_y x + (1+a)b_x y - b_x b_y = 0. \tag{2.3}$$

The proof is given in appendix A. Note: the affine transformation has only three degrees of freedom (given by a, b_x, b_y in the statement of the theorem); and, the theorem also applies if the original object has affine skewed symmetry, since an affine image of an affine transformed object with bilateral symmetry has the same properties.

(*b*) *Back-projection*

In this section we consider the extent to which we can unskew images such as those shown in figure 1. Evidently, some of the skewed symmetries in figure 1 have only a single symmetry, in which case the best one can hope for in general (that is, without mobilizing further knowledge) is to unskew the image to a single-parameter family of symmetric shapes that corresponds to tilting the symmetric shape backwards while preserving the direction of the symmetry axis. If, on the other hand, there is more than one symmetry axis, then one might hope to combine the information from two or more such axes to uniquely unskew the shape. This insight is embedded in the following theorem.

Theorem 2.2. *Suppose we have an (uncalibrated) image of one or more coplanar symmetric objects.*

One symmetry. If there is only one symmetry present, the image can be back-projected, modulo a similarity, to form a one-parameter family of symmetric shapes that could have given rise to the image.

Two symmetries. In the case that two symmetries are present in the image, the image can be back-projected uniquely, modulo a similarity, provided that the two symmetry axes are neither parallel nor orthogonal (either in the image or in space, since the camera is assumed affine).

Proof. First we give some notation for back-projections. Suppose the affine transformation relating the object and image planes is given by

$$\boldsymbol{X} = U\boldsymbol{x} + \boldsymbol{B}, \tag{2.4}$$

where $\boldsymbol{x}$ is the two-dimensional image point, $\boldsymbol{X}$ is the corresponding point in the back-projected planar scene, $\boldsymbol{B}$ is a two-vector translation and U is the 2×2 linear transformation matrix with $\det U \geqslant 0$ responsible for back-projection. It can be shown (Blake & Marinos 1990) that the transformation U can be decomposed as

$$U = \lambda_1 R(\theta) P(\lambda, \tau), \tag{2.5}$$

where $R(\theta)$ is a rotation by θ, and $P(\lambda, \tau)$ is a symmetric matrix:

$$P(\lambda, \tau) = R(\tau) \begin{bmatrix} \lambda & 0 \\ 0 & 1 \end{bmatrix} R(-\tau), \tag{2.6}$$

with

$$R(\tau) = \begin{bmatrix} \cos\tau & -\sin\tau \\ \sin\tau & \cos\tau \end{bmatrix}.$$

This particular decomposition of U makes its four degrees of freedom explicit in a way that corresponds to the processes of projection: the linear transformation U consists of an isotropic scaling by λ_1, a rotation about the optic axis by $R(\theta)$, and an expansion by λ in the direction $\boldsymbol{\tau} = (\cos\tau, \sin\tau)$ (the eigenvector of $P(\lambda, \tau)$).

The eigenvectors of A, say $\boldsymbol{a}$ and $\boldsymbol{b}$, back-project to vectors parallel to and orthogonal to the symmetry axis, respectively. In the object plane, therefore, the scalar product $(U\boldsymbol{a}) \cdot (U\boldsymbol{b}) = 0$ and, consequently,

$$\boldsymbol{a}^\top U^\top U \boldsymbol{b} = 0. \tag{2.7}$$

The matrix $V = U^\top U$ is symmetric and positive definite. Let the components of V be given by

$$V = \begin{bmatrix} \alpha & \beta \\ \beta & \gamma \end{bmatrix}, \tag{2.8}$$

then we have from equation (2.7)

$$\begin{pmatrix} a_x b_x & a_x b_y + a_y b_x & a_y b_y \end{pmatrix} \begin{bmatrix} \alpha \\ \beta \\ \gamma \end{bmatrix} = 0. \tag{2.9}$$

This is a linear constraint on α, β, γ. Two such constraints determine the *ratio* $\alpha : \beta : \gamma$. The sign is fixed by the requirement that V is positive definite, so that $\operatorname{tr} V = \alpha + \gamma > 0$. This is sufficient to determine λ and τ, as is shown in the following lemma.

Lemma 2.3. *The ratio $\alpha : \beta : \gamma$, with sign chosen so that $\alpha + \gamma > 0$, determines λ up to a four-fold ambiguity, and τ up to a two-fold ambiguity.*

Note that λ^2 is the ratio of the eigenvalues of V, and τ is the rotation angle to the eigendirections. From equations (2.5) and (2.6),

$$V = \lambda_1^2 P(\lambda^2, \tau). \tag{2.10}$$

The trace and determinant of V give two equations for λ^2:

$$\operatorname{tr} V = \alpha + \gamma = \lambda_1^2 (1 + \lambda^2),$$
$$\det V = \alpha\gamma - \beta^2 = \lambda_1^4 \lambda^2.$$

Eliminating λ_1^2 gives

$$\frac{(\operatorname{tr} V)^2}{(\det V)} = \frac{(\alpha + \gamma)^2}{\alpha\gamma - \beta^2} = \frac{(1 + \lambda^2)^2}{\lambda^2} = 4\mu. \tag{2.11}$$

Solving this gives $\lambda^2 = \{\lambda'^2, 1/\lambda'^2\}$ where $\lambda'^2 = 2\mu - 1 + 2\sqrt{\mu(\mu - 1)}$. Or, equivalently, $\lambda = \pm\sqrt{\mu} \pm \sqrt{\mu - 1}$, which are four solutions of the form $\{\lambda', 1/\lambda', -\lambda', -1/\lambda'\}$ with $\lambda' = \sqrt{\mu} + \sqrt{\mu - 1}$.

The rotation angle τ is obtained from equations (2.6) and (2.10):

$$\begin{bmatrix} \alpha & \beta \\ \beta & \gamma \end{bmatrix} = \lambda_1^2 R(\tau) \begin{bmatrix} \lambda^2 & 0 \\ 0 & 1 \end{bmatrix} R(-\tau). \tag{2.12}$$

Rearranging gives

$$\begin{bmatrix} \cos\tau & \sin\tau \\ -\sin\tau & \cos\tau \end{bmatrix} \begin{bmatrix} \alpha & \beta \\ \beta & \gamma \end{bmatrix} \begin{bmatrix} \cos\tau & -\sin\tau \\ \sin\tau & \cos\tau \end{bmatrix} = \lambda_1^2 \begin{bmatrix} \lambda^2 & 0 \\ 0 & 1 \end{bmatrix}.$$

From the off-diagonal elements we obtain

$$-\sin^2\tau\beta - \cos\tau\sin\tau(\alpha - \gamma) + \cos^2\tau\beta = 0,$$

from which

$$\tan 2\tau = \frac{2\beta}{\alpha - \gamma}. \tag{2.13}$$

This gives four solutions for τ. If $\tau^* = \frac{1}{2}\arctan[(2\beta)/(\alpha - \gamma)]$, then the four solutions are $\{\tau^* + n\pi/2\}$ for $n = 0, 1, 2, 3$. Two of the solutions are simply due to a clockwise rotation, as opposed to counter-clockwise, so may be disregarded. It is only necessary to determine whether τ is in the first or second quadrant. From (2.12),

$$\beta = \lambda_1^2(\lambda^2 - 1)\cos\tau\sin\tau = \tfrac{1}{2}\lambda_1^2(\lambda^2 - 1)\sin 2\tau,$$

so that $\operatorname{sgn}(\beta/(\lambda^2 - 1)) = \operatorname{sgn}(\sin 2\tau)$, and if $\operatorname{sgn}(\beta/(\lambda^2 - 1)) \geqslant 0$ then $0 \leqslant \tau \leqslant \pi/2$, otherwise $\pi/2 < \tau \leqslant \pi$. Consequently, there are two solutions for τ corresponding to the two solutions $\{\lambda'^2, 1/\lambda'^2\}$ above for λ^2. ■

This proves the lemma. Now we can return to the proof of the theorem. In the following we take $\lambda = \sqrt{\mu} + \sqrt{\mu - 1}$. The other solutions differ only by similarity transformations. If $\lambda = \sqrt{\mu} + \sqrt{\mu - 1}$, then for real solutions $\mu \geqslant 1$, and, consequently, $\lambda \geqslant 1$. Hence, $\operatorname{sgn}(\beta) = \operatorname{sgn}(\sin 2\tau)$ and this uniquely determines τ.

One symmetry. Equation (2.9) has a one-parameter family of solutions for the ratio $\alpha : \beta : \gamma$. Correspondingly, there is a one-parameter family of solutions for λ and τ.

Two symmetries. Two symmetries generate two constraint equations (2.9):

$$M \begin{bmatrix} \alpha \\ \beta \\ \gamma \end{bmatrix} = 0,$$

where

$$M = \begin{bmatrix} a_x^1 b_x^1 & a_x^1 b_y^1 + a_y^1 b_x^1 & a_y^1 b_y^1 \\ a_x^2 b_x^2 & a_x^2 b_y^2 + a_y^2 b_x^2 & a_y^2 b_y^2 \end{bmatrix}. \tag{2.14}$$

Provided the matrix M is of rank two, this uniquely determines the ratio $\{\alpha : \beta : \gamma\}$ (and, consequently, from the lemma, λ and τ). It can be shown that M drops rank if any of the vectors $\{\boldsymbol{a}^1, \boldsymbol{b}^1, \boldsymbol{a}^2, \boldsymbol{b}^2\}$ are parallel, hence the clause in the theorem. The optimal solution when there are more than two constraints is discussed in § 4. ■

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

3. Detecting symmetries

As shown in § 2, corresponding sides Γ and Γ' of a symmetric planar object project to image curves γ and γ', respectively. Even though the projection ($\Gamma \rightarrow \gamma$ and $\Gamma' \rightarrow \gamma'$) is by a general affine transformation, the image curves γ and γ' are related by a three-dimensional subset of the (six degrees of freedom) planar affine group. In this section we describe how these results can be utilized to detect such image pairs efficiently.

Intra-image curve matching has much in common with the inter-image curve matching necessary for model-based vision, and approaches developed for that area can be used to advantage here. In particular, and this introduces the second theme of invariants foreshadowed in the introduction, the use of invariants as index functions avoids the cost of a six-dimensional search over transformation parameters (Lamdan *et al.* 1988; Rothwell *et al.* 1992*c*). The three stages of an implemented algorithm are described in the following sections.

(*a*) *Generating and matching affine invariants*

Two curves that are related by an affine transformation have the same affine invariants. The converse is not necessarily true, but invariants can usefully be used to generate hypotheses for matching, which can subsequently be tested. Briefly, a function $I(\Gamma)$ of a curve Γ is an *invariant* if $I(\gamma) = |U|^w I(\Gamma)$, where γ is the image of Γ (refer to figure 2) under the affine transformation as defined in equation (2.4). The exponent w is the weight of the invariant. If $w = 0$ then the invariant is absolute, otherwise it is relative. Note that in order to determine *local* symmetries, the invariant must not depend on global properties of the curve. Examples of (semi-local) affine invariants for smooth curves are given below.

Unlike model-based vision, where *absolute* invariants are needed, *relative* invariants suffice in this case. To see this, consider two symmetry-related curves in the object plane. These are related by an affine transformation with determinant -1 (since the transformation is a reflection). Consequently, affine invariants of each side of the shape are equal modulo a sign. In the image, invariants are multiplied by $|U|^w$, which is unknown, but which is the same for both sides. Thus, relative affine invariants of each side have the same magnitude.

Matching on invariants can be implemented as an $O(n)$ complexity process by the use of hashing (where n is the number of curves) (Rothwell *et al.* 1992*a*, *c*). We have implemented the simpler $O(n^2)$ algorithm, since n is small in the cases we have experimented with. It is straightforward to implement the more complicated algorithm.

(*b*) *Determining the affine transformation*

Having found two curves with matching invariant(s), the next stage is to determine if the curves are affine related. This is achieved by extracting a number of *distinguished points* on each curve, and determining an affine transformation between these point sets. Three points are required to determine a general affine transformation. Distinguished points are curve 'markers' that can be determined before and after a transformation. A number of examples are shown in figures 3–5. They include points preserved by projectivities (such as inflections, bitangent contact points, 'cast' tangents) as well as those exclusively preserved by affinities (such as points defined by parallel lines). Note, points are ordered by the curve, so the correspondence problem is greatly simplified.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

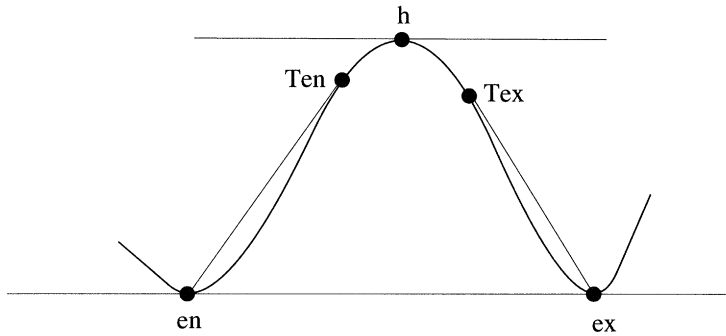

Figure 3. Examples of distinguished points for a non-convex curve under affine transformations. Points en and ex mark the entrance and exit of the concavity determined by the bitangent line. Further distinguished points are constructed from these points: Ten is the point on the curve which is tangent to a ray based at en (similarly for Tex and ex); h is determined by the line parallel and furthest from the bitangent. Apart from h, these distinguished points are also preserved by projective transformations. Examples of these points on an image curve are shown in figure 5.

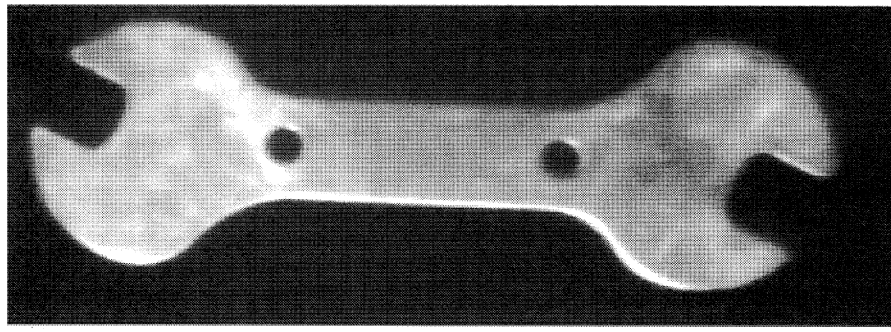

Figure 4. Affine view of spanner.

(*c*) *Verifying subset membership*

As noted in § 2, the image curves are related not by a general affine transformation (with six degrees of freedom), but by a three-parameter subspace. If the affine transformation does not lie in this subspace, then the two curves cannot be symmetry related. Note that if two curves *are* symmetry related then two points are sufficient to determine the transformation of equation (2.2). When more points are available, the form of the transformation is used as a constraint (via a Lagrange multiplier) in a least-squares estimator. Details are given in § 3 *e*. If the curves are affine related then one can be superimposed onto the other by applying the computed affine transformation. An example is shown in figure 10.

(*d*) *Affine semi-local invariants*

For a non-convex curve, following Lamdan *et al.* (1988), we exploit concavities, by constructing a bitangent across the concavity and determining the interior point on the concavity curve with tangent parallel to the bitangent. See figure 3. This generates three distinguished points. This particular choice of points has the advantage that it does not depend globally on the curve. Consequently, if part

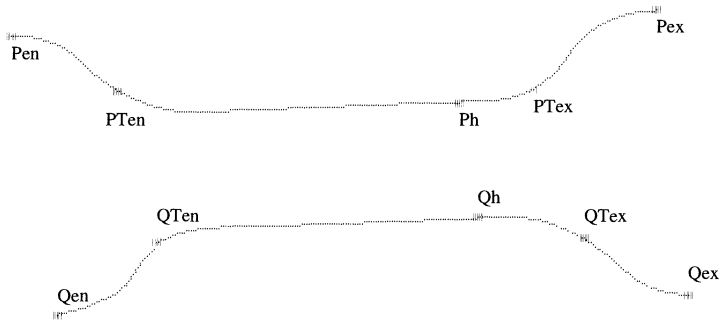

Figure 5. Distinguished points for edge curves extracted from the image of figure 4. The notation for the curves P and Q is defined in figure 3.

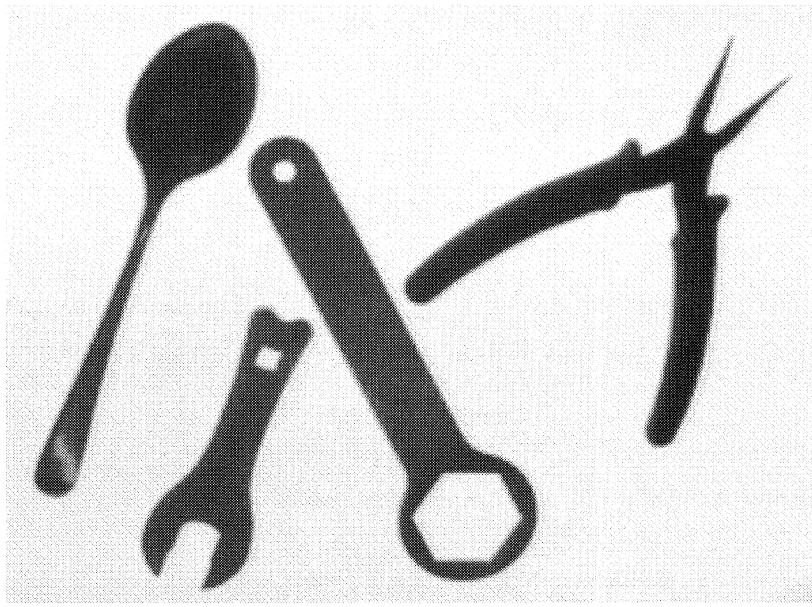

Figure 6. Image of four objects with bilateral symmetries.

of the curve is occluded or missed because of segmentation problems, *local* symmetries can still be detected. Affine invariants are generated from the concavity curve.

Area in the image space. This is a relative invariant. The area used is that of the triangle defined by the three concavity distinguished points (en, ex and h).

Moments in the canonical frame. Significant concavities are mapped to a canonical frame (Lamdan *et al.* 1988) consisting of an equilateral triangle with vertices at $(-1, 0)$, $(1, 0)$ and $(0, \sqrt{3})$ by using the affine-basis triplet points of the concavities. The x and y moments of the concavity in the canonical frame are used as invariant indexes.

Table 1 lists the invariant values for the objects in figure 6. These differ, in general, by less than 2% for symmetry-related concavities. Figure 7 shows the matched concavity pair extracted from figure 6. For the pliers, invariant values

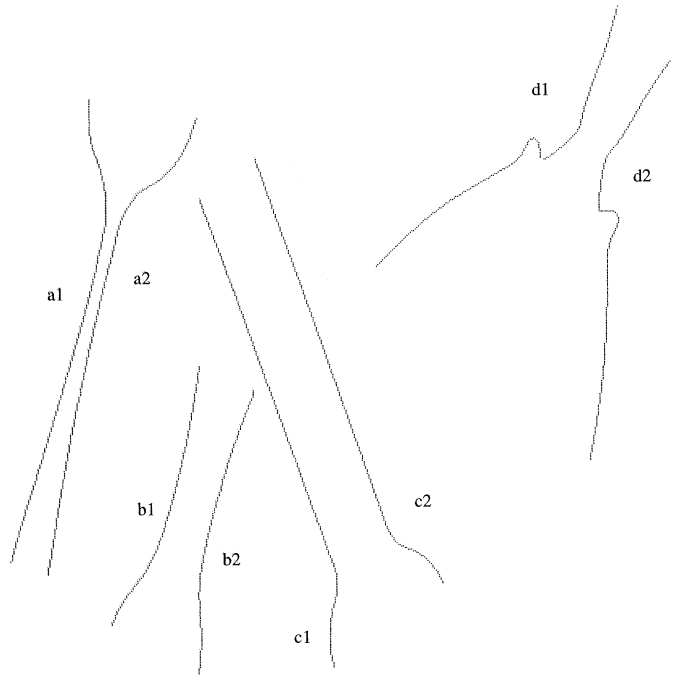

Figure 7. Matched concavities (a1, a2), (b1, b2), (c1, c2) and (d1, d2) are extracted from the symmetric objects of figure 6.

Table 1. *Affine invariant values for objects in figure 6*

invariant	spanner		spoon		hex spanner		pliers	
area (image space)	1064	1056	3694	3855	2719	2723	4398	3814
moment about x-axis (canonical frame)	10.85	11.38	32.13	31.44	18.61	17.13	2.57	2.03
moment about y-axis (canonical frame)	102.15	99.71	86.17	84.41	92.72	93.83	82.59	78.82

are not consistent because of the thickness of the handles. The handles cause two problems: first, they are rounded so that (as in the case of an extremal boundary), in general, the surface curves projecting to the outline will be space curves and not mirror pairs; and second, and more important in this case, the handles and jaw are not in the same plane.

For a convex curve segment, distinguished points can be obtained from the anti-symmetry set (Blake & Taylor 1993).

(*e*) *Implementation and results*

Feature extraction. Image contours are extracted using a local implementation of the Canny (1986) edge detector. Significant concavities are extracted for each closed contour after computing a convex hull and setting a threshold on concavity

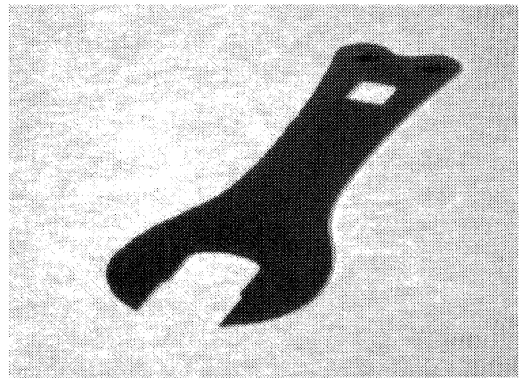

Figure 8. Affine view of symmetric spanner.

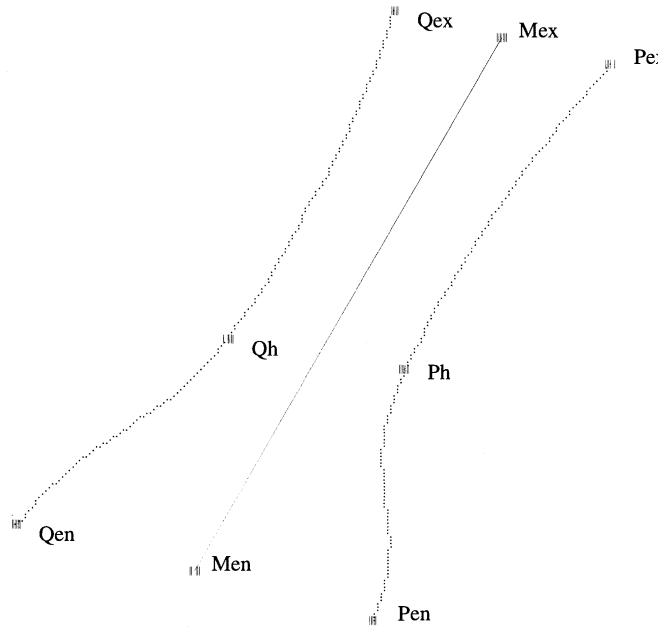

Figure 9. Matched concavities extracted from symmetric spanner of figure 8.

height and width. Bitangents are found via a dual space construction (Rothwell *et al.* 1992*a*) and determine the concavity entrance (en) and exit points (ex).

Concavity matching. For each closed contour in the scene, matched concavities are detected using affine invariant indexes as described in § 3 *a*. Examples of the three points, en, ex and h, used as an affine basis are shown in figure 9. Corresponding points in the matched concavity pair are determined from the tracing order (clockwise or anti-clockwise) of the image contour from which concavities are extracted. These point correspondences are used to determine the affine transformation between corresponding curves.

Affine transform. The next step is to determine if the affine transformation arises from a reflectional symmetry of the object curves, i.e. whether it lies in the three-degrees-of-freedom subspace defined by the constraints of equation (2.2). In

Table 2. *As the accuracy of the affine transformation is improved it more closely satisfies the constraints of equation (2.2)*

(The table gives the computed A and $\boldsymbol{b}$ elements when they are calculated either: (i) directly from three points; or (ii) via a pseudo-inverse from five points. The transformation is for the points obtained from figure 21.)

number of points	affine elements solving equations (2.1) and (3.2)						elements calculated as equation (2.2)	
	a_{11}	a_{12}	b_x	a_{21}	a_{22}	b_y	$a_{12} = -b_x(1+a)/b_y$	$a_{21} = -b_y(1-a)/b_x$
3	0.89	0.63	−72.16	0.49	−0.79	158.57	0.86	0.24
5	0.77	0.97	−87.79	0.43	−0.75	157.15	0.98	0.41

practice, the affine transformation determined from the three affine basis points is not sufficiently accurate, so extra correspondences are included via a pseudo-inverse. Two additional points, marked Ten, Tex in figures 3 and 5, are the points of tangency to the extracted curve drawn through the cavity entrance and exit points (these are determined from the convex hull). In straightforward notation, equation (2.1) is rewritten as

$$P\boldsymbol{X} = Q, \tag{3.1}$$

where $\boldsymbol{X}$ is a vector formed from the elements of A and $\boldsymbol{b}$. This is solved using a pseudo-inverse to give

$$\boldsymbol{X} = P^{\top}(PP^{\top})^{-1}Q. \tag{3.2}$$

Table 2 demonstrates that as the number of points increases, the accurately determined affine transformation does indeed satisfy the constraints of equation (2.2).

Having determined the affine transformation, the symmetry axis is given by (2.3). Since midpoints are preserved by affine transformations, the midpoints of lines joining corresponding distinguished points lie on the symmetry axis, and this provides a quick, though not as accurate, method for determining the line.

An alternative method for improving the accuracy of the affine transformation is to minimize differences between the curve on one side and the other affine-transformed side (so it should be identical). For example, differences of area or the distance between corresponding points of the matched curve could be used as a measure. This has not been implemented, but the accuracy of the affine transformation computed from the pseudo-inverse is demonstrated in figure 10, where one side Q is 'reflected' onto the other side P.

(*f*) *Global symmetries*

Clearly, a local symmetry between two concavities does not imply a global reflectional symmetry for the whole object. For example, while the symmetric spanner shown in figure 8 has a global reflectional symmetry, the spanner shown in figure 4 does not. To test for global symmetry, the local symmetry line is extended in both directions, while there is evidence that a symmetry with this axis exists. This is the case if for each point on one side there is a corresponding point on the other side in the direction of $\boldsymbol{b}$ (as defined in equation (2.2)) at the same distance from the symmetry line.

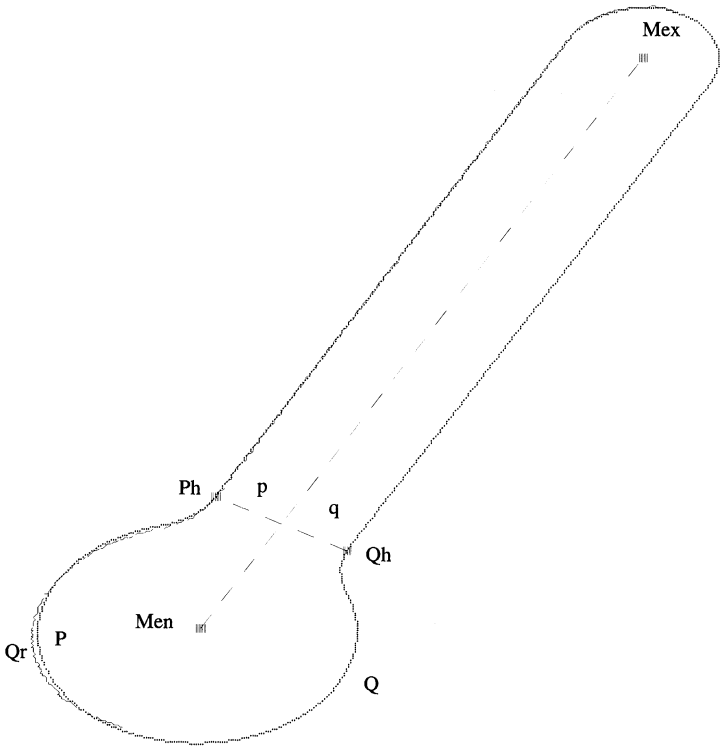

Figure 10. Qr is the image 'reflection' of Q onto P using an affine transformation computed from five point correspondences for the hex spanner shown in figure 21.

Figure 11 is an affine scene containing globally, partially symmetric and non-symmetric objects. Globally symmetric contours are correctly determined as shown in figure 12.

4. Applications

(*a*) *Back-projection*

Here we determine the affine back-projection to the object plane using the results of § 2. Note, back-projection does not require camera aspect ratio (or any of the intrinsic parameters). We first give an intuitive and simple construction for determining the back-projection and its uniqueness, which is applicable for up to two symmetries.

Consider an image consisting of two coplanar objects with single bilateral symmetries. Determine the skewed symmetry axes of each object (say, by joining the midpoints of corresponding distinguished points), and choose an origin on one of the symmetry axes, with vector $\boldsymbol{a}$ on the axis, and $\boldsymbol{b}$ parallel to the lines joining corresponding points. See figure 13. The back-projection is achieved in two stages.

Stage one. Unskew the first object by determining the transformation that maps $\boldsymbol{a}$ and $\boldsymbol{b}$ to the points $(0, 1)$ and $(\nu, 0)$ (*unskewing frame*). This determines

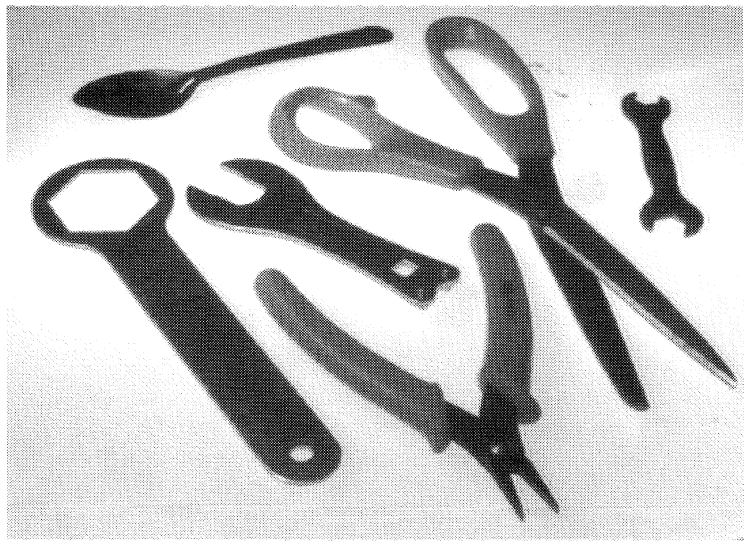

Figure 11. Affine scene of globally, partially and non-symmetric objects.

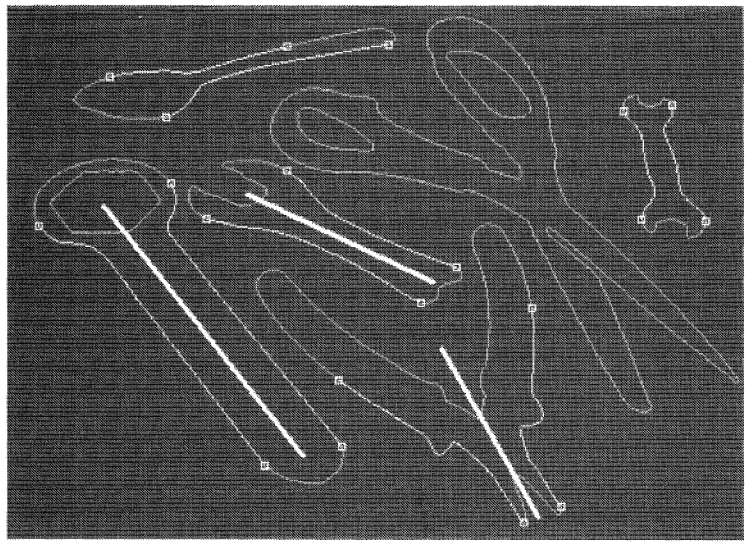

Figure 12. Outline curves, concavity entrance and exit point for the objects in figure 11. Symmetry lines are only drawn where the object is determined to be globally symmetric.

three of the degrees of freedom of U, including the arbitrary rotation and isotropic scaling, but does not determine the object-plane aspect ratio. Explicitly,

$$\begin{bmatrix} u_{11} & u_{12} \\ u_{21} & u_{22} \end{bmatrix} \begin{bmatrix} a_x^1 & b_x^1 \\ a_y^1 & b_y^1 \end{bmatrix} = \begin{bmatrix} 0 & \nu \\ 1 & 0 \end{bmatrix}.$$

So,

$$U = \begin{bmatrix} 0 & \nu \\ 1 & 0 \end{bmatrix} \begin{bmatrix} a_x^1 & b_x^1 \\ a_y^1 & b_y^1 \end{bmatrix}^{-1},$$

yielding the expected one-parameter (ν) family of solutions.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

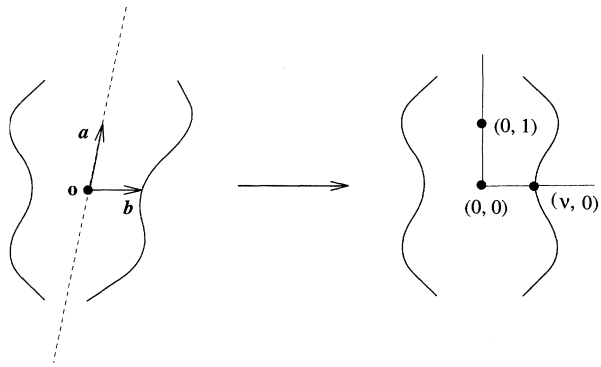

Figure 13. Mapping of guiding vectors $\boldsymbol{a}$ and $\boldsymbol{b}$ to unskewing frame.

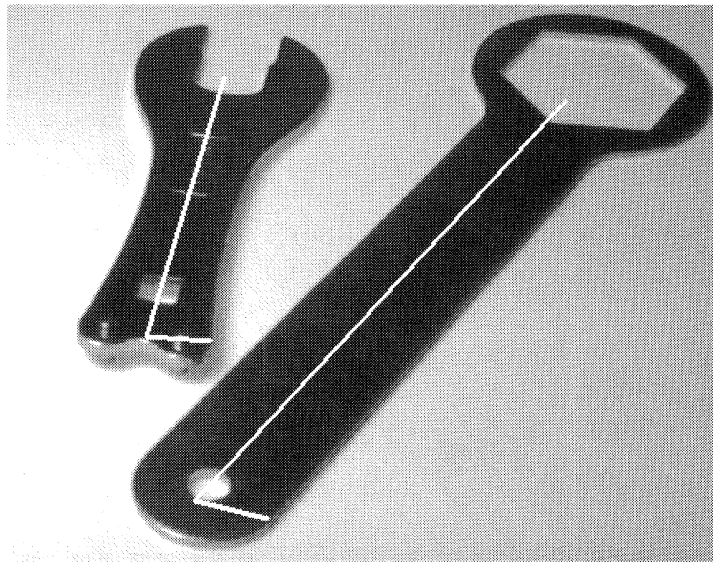

Figure 14. Affine view of two objects with superimposed guiding vectors. The vectors are clearly not perpendicular.

Stage two. Now, ν is determined by enforcing that the second object should also be unskewed. We have:

$$\boldsymbol{a}^{2t} U^{\top} U \boldsymbol{b}^2 = 0. \tag{4.1}$$

Multiplying out gives:

$$\boldsymbol{a}^{2t} \begin{bmatrix} a_x^1 & b_x^1 \\ a_y^1 & b_y^1 \end{bmatrix}^{-t} \begin{bmatrix} 1 & 0 \\ 0 & \nu^2 \end{bmatrix} \begin{bmatrix} a_x^1 & b_x^1 \\ a_y^1 & b_y^1 \end{bmatrix}^{-1} \boldsymbol{b}^2 = 0, \tag{4.2}$$

which is a linear equation for ν^2. Note that if $\boldsymbol{b}^1$ is parallel to $\boldsymbol{b}^2$ (and consequently $\boldsymbol{a}^1$ is parallel to $\boldsymbol{a}^2$) then the quadratic form in (4.2) is identically zero and there is no constraint on ν. Similarly, there is no constraint if $\boldsymbol{a}^1$ is parallel to $\boldsymbol{b}^2$ (and consequently $\boldsymbol{b}^1$ is parallel to $\boldsymbol{a}^2$). This occurs if the symmetry axes of both objects are parallel or orthogonal. In this case both objects are unskewed by the first stage.

This formulation is, of course, equivalent to § 2 *b*, and either can be used if there

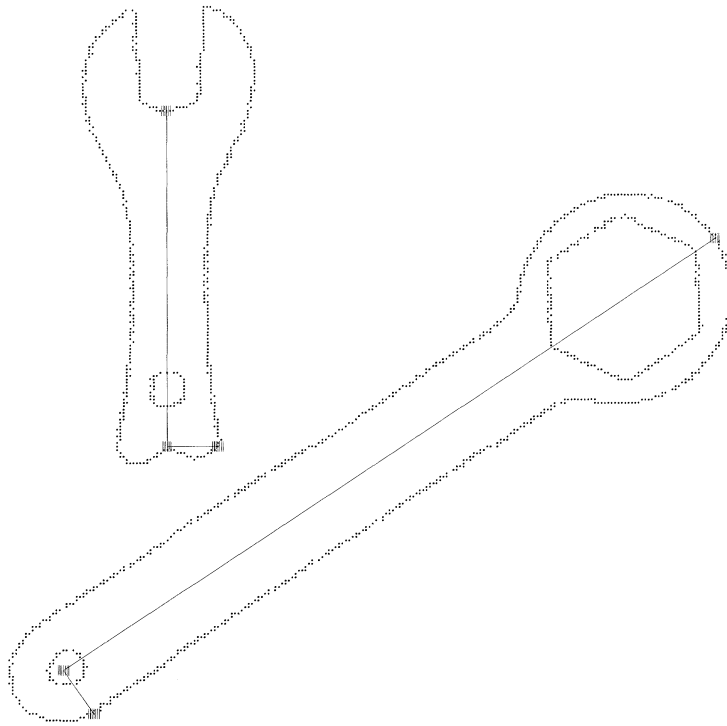

Figure 15. Unskewed image of the spanner and hex spanner of figure 14.

are two symmetries present. If there are more than two symmetries, where a least-squared solution is required, then the above method is not easily generalizable. However, the formulation of § 2 *b* is not restricted. Its application in a least-squared solution is described below.

Figures 14 and 15 show examples of object pairs before and after back-projection. The angle between the vectors $\boldsymbol{a}^i$ and $\boldsymbol{b}^i$ before and after are given in table 3. Accurate back-projection requires accurate determination of these vectors (which are the eigenvectors of the matrix A). In practice, we find that five point correspondences and use of a pseudo-inverse, as described in § 3 *e*, are sufficient to determine A to a satisfactory accuracy.

Up to this point it has been assumed that each object only contributes a single local symmetry. However, should an object contain several local symmetries then this object alone is sufficient to determine the back-projection (provided the usual conditions are satisfied). An example is shown in figure 16. We have taken two local symmetries enforcing the constraint that they cannot be mutually parallel or perpendicular. Local symmetries with guiding vectors and the matching concavities are shown in figure 17. Figure 18 shows the unskewed image of the object and table 4 demonstrates the back-projection. Slant and tilt are also determined in this case. This is described in § 4 *b*.

Least-squares solution. If there are more than two coplanar objects, then the transformation is estimated using a least-squares technique. Specifically, equation (2.9) is a linear constraint on the kernel vector (α, β, γ). If there are n

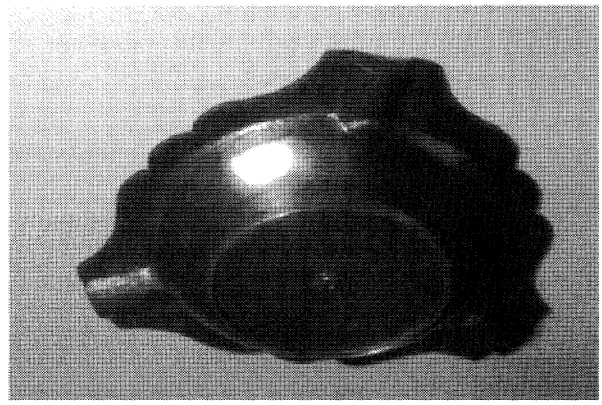

Figure 16. Affine image of an object with multiple local symmetries.

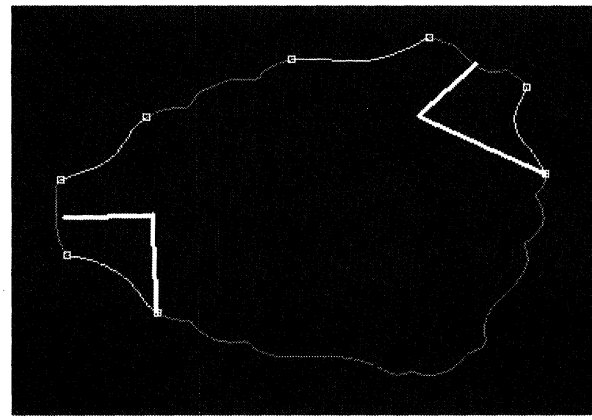

Figure 17. Local symmetries and matching concavities extracted from figure 16. Note that local symmetries considered are neither parallel nor perpendicular.

Table 3. *Angles between guiding-point triplet before and after unskewing for objects in figure 14*

objects	spanner	hex spanner
initial angle	76.4	81.9
final angle	90.0	89.9

objects, we seek the minimum of $\|M\boldsymbol{x}\|^2$ subject to $\|\boldsymbol{x}\| = 1$, where M is a $n \times 3$ matrix with each row given by equation (2.9). This is a standard problem in linear algebra. The solution is the eigenvector of $M^\top M$ with least eigenvalue. It is also possible to determine a covariance matrix for λ and μ in a similar manner to Blake & Marinos (1990). A more complete treatment of image noise and segmentation errors would weight each row of M according to a measure of its uncertainty.

Results of applying this least-squares estimator to the affine scene in figure 19

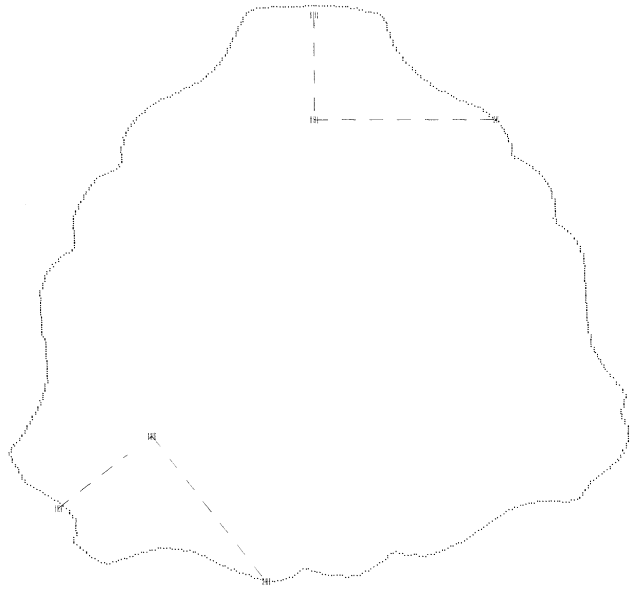

Figure 18. Unskewed image of the object in figure 16.

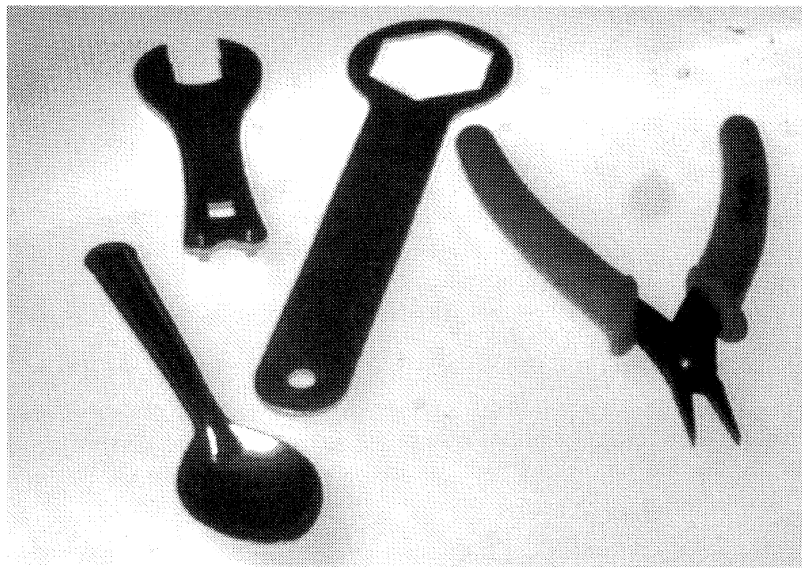

Figure 19. Affine scene of multiple objects.

are given in table 5 and figure 20. Note that the angle between the guiding vectors for the pliers is not as good as the others due to the handle limitation discussed in § 3 *d*.

Back-projection is generally formulated as maximizing a function – in this case one sensitive to the angle between back-projected guiding vectors, but unaffected by similarity transformations. Instead of the linear constraint given in equation (2.9), back-projection could be computed by minimizing the nonlinear

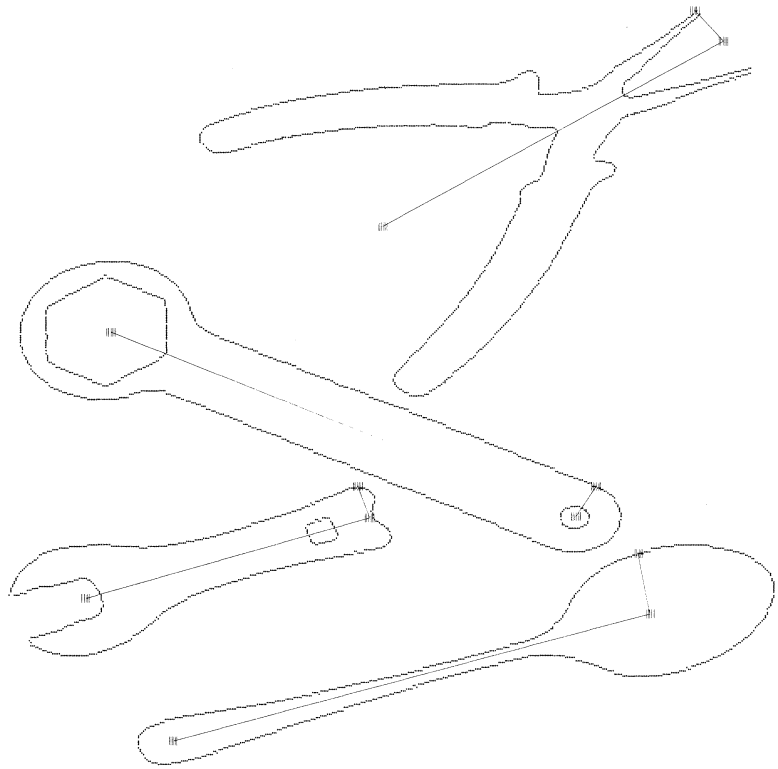

Figure 20. Unskewed image of multiple objects of figure 19.

Table 4. *Angles between guiding vectors before and after unskewing the object in figure 16*

objects	first symmetry	second symmetry
initial angle	95.8	74.3
final angle	90.0	90.0

function $f(x) = \sum \cos^2 \theta_i$ (where θ_i is the skew angle between $\boldsymbol{a}^i$ and $\boldsymbol{b}^i$), and the sum includes all symmetric objects in the scene.

Combination with other constraints. If there is only one symmetry in the image, this is only sufficient to determine the back-projection (modulo similarity) up to a one-parameter family. However, other scene-specific information can be incorporated to resolve the ambiguity.

A number of constraints can be put forward in the case of the hex spanner (figure 21). For example, after unskewing, all the sides of the hexagonal head should have equal length and the conic surrounding the hexagonal head should be circular.

More generally, for a regular isotropic shape, like a hexagon in this case, compactness, ((area)/(perimeter)^2) as proposed by Brady & Yuille (1984) and Horaud

Table 5. *Angles between guiding point triplet before and after unskewing for multiple objects in figure 19*

objects	spanner	spoon	hex spanner	pliers
initial angle	76.4	97.2	81.5	94.3
final angle	89.7	89.1	89.6	88.8

Figure 21. Skewed image of hex spanner with superimposed guiding vectors.

& Brady (1988), is maximized in the object frame. The result of removing the ambiguity by imposing equality of the hexagonal sides is shown in figure 22.

(*b*) *Slant and tilt determination*

If the camera aspect ratio is known, so that camera projection is scaled orthography rather than affine, then the slant and tilt of the object plane can be determined from the ratio $\{\alpha : \beta : \gamma\}$. Quite straightforwardly, the variables λ and τ which appear in back-projection operator (2.6) are, respectively, $\sec\sigma$, where σ is the slant of the object plane, and tilt (with $\lambda^2 = \lambda'^2$ so that $|\lambda| > 1$). See figure 23 and Blake & Marinos (1990).

The ratio $\{\alpha : \beta : \gamma\}$ determines λ up to sign, corresponding to a reflection of the plane. Thus, slant is recovered up to the usual two-fold ambiguity under scaled orthographic projection, i.e. σ and $\pi - \sigma$.

Calculated slant and tilts are given in table 6. The results are compared to: (i) slant and tilt obtained by a method which back-projects a circle under perspective (Rothwell *et al.* 1992*b*); and (ii) approximate measurements from the camera position. Two camera orientations are compared. Three results are given for each orientation corresponding to different arrangements of the coplanar objects. One representative image for each orientation is shown in figure 24.

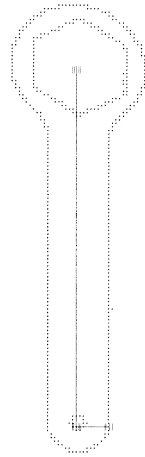

Figure 22. Unskewed image of hex spanner of figure 21.

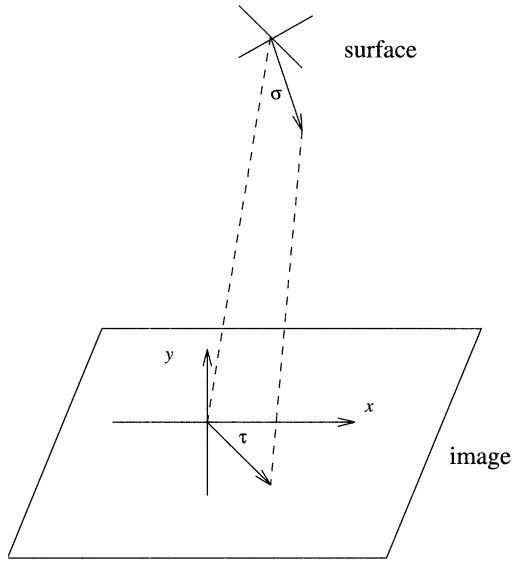

Figure 23. Slant (σ) and tilt (τ) angle.

(*c*) *Planarity tests*

Suppose two symmetric planar objects are not coplanar; can this be detected from the image? If the objects are not coplanar then equations (2.11) or (4.1) may not have a solution. This provides a simple test for non-coplanarity which will always be passed if the objects *are* coplanar (subject to image noise), but which non-coplanar objects may fail.

The test is derived from equation (2.11) as follows. We have

$$\frac{(1+\lambda^2)^2}{\lambda^2} = \frac{(\alpha+\gamma)^2}{\alpha\gamma-\beta^2} = 4\mu. \tag{4.3}$$

Considering the λ equality first, a number of constraints may be evaluated as follows:

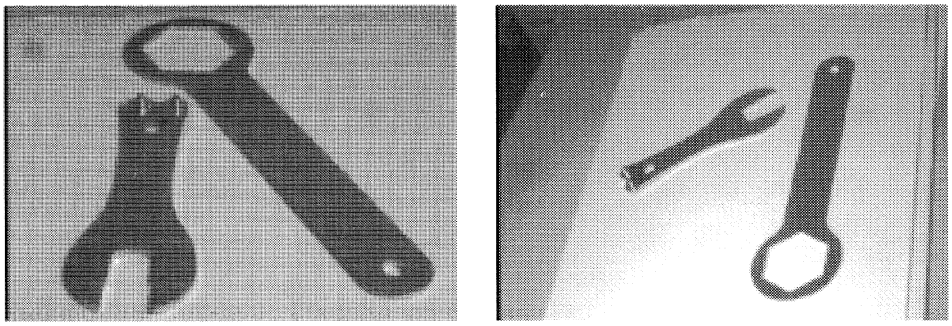

Figure 24. Examples of images used to calculate σ and τ for each of the camera orientations. For each orientation slant and tilt is recovered for three different object arrangements. Results are given in table 6.

Table 6. *Slant and tilt values calculated from symmetry back projection, circular back-projection and (approximate) camera position*

orientation	symmetry		circle		measured
	σ	τ	σ	τ	σ
figure 24 left	50 ± 3	88 ± 2	53	90	54
figure 24 right	41 ± 3	106 ± 4	47	118	44
figure 16	44	111	47	118	44

(i) for λ^2 to be real and positive, $\mu \geqslant 0$ (from $\mu \geqslant 0$, it follows that $\alpha\gamma \geqslant \beta^2$ or $\det V \geqslant 0$, where V is defined in equation (2.8)); and

(ii) multiplying out gives $\lambda^2 \pm 2\sqrt{\mu}\lambda + 1 = 0$. This only has real roots if the discriminant $\mu - 1 \geqslant 0$, i.e. for λ real, $\mu \geqslant 1$.

This defines a region for acceptable $\{\alpha : \beta : \gamma\}$ solutions. If solutions do not lie in this region then the image cannot have arisen from coplanar objects.

An example is shown in figure 25, where the hex spanner is in a different plane from the other spanner. The calculated values of $\{\alpha : \beta : \gamma\}$ and μ are $\{-0.793 : -0.0920 : 1\}$ and -0.013, respectively. The above planarity-test condition demonstrates that the objects are not in the same plane.

5. Discussion

We have demonstrated that the object relation of bilateral symmetry gives rise to image constraints that can be utilized in real applications. In particular, bilateral symmetry restricts the affine transformation between corresponding image contours to a three-dimensional subset of the planar affine group. This constraint allows these contours to be discriminated from other affinity-related image curves. Similar constraints can be derived for other object relations. For example, a planar object with two-fold rotational symmetry induces the following constraints on the affinity relating corresponding image curves:

(i) $A^2 = I$;

(ii) $|A + I| = 0$; and

(iii) $|A - I| \neq 0$.

Figure 25. Affine image of two non-coplanar objects.

and similarly for n-fold rotational symmetry (with $A^n = I$).

As well as extending to other relations, the approach can be extended to the most general planar object in image transformation – a projective transformation. This is described in appendix B. Unlike the affine case, lines joining corresponding image points are not, in general, parallel. However, the symmetry restricts the transformation between corresponding image curves to a subset of the full group – here to a four-dimensional subset of the plane projective group (an involution). This has an application in the recognition of objects which are surfaces of revolution, since the two 'sides' of the apparent contour of these objects are projectively equivalent to curves with bilateral symmetry. Semi-local invariants can again be used to indicate matches, although absolute, not relative, invariants must be used in the projective case.

Another extension is to affine images of three-dimensional objects with bilateral symmetry. Exploiting this constraint facilitates the recovery of three-dimensional structure and pose from single images. It can be shown that structure can be recovered, modulo a Euclidean transformation, to a four-parameter family of symmetric objects that could have given rise to the image (Fawcett *et al.* 1993). If the object has two orthogonal bilateral symmetries, the shape can be reconstructed modulo similarity.

We are very grateful for helpful comments from Andrew Blake, Ron Daniel, David Forsyth, D. Dutta Majumder, Joe Mundy and particularly Charlie Rothwell and the anonymous referees. D.P.M. acknowledges the support of the United Nations Development Project fellowship awarded to the KBCS Project, Indian Statistical Institute, Calcutta, India. Financial support was provided by ESPRIT Project 6448 'VIVA', and by the Science and Engineering Research Council.

Appendix A. Proof of theorem 2.1

Suppose two image curves γ and γ' are images of two sides of a planar object with bilateral symmetry, and image projection can be represented by an affine transformation. Then the transformation between γ and γ' has the following properties.

Property 1. γ and γ' are related by an affine transformation. That is, if $\boldsymbol{x}$ is a point on γ then there is a point $\boldsymbol{x}'$ on γ' such that

$$\boldsymbol{x}' = A\boldsymbol{x} + \boldsymbol{b},$$

where A is a non-singular 2×2 matrix, and $\boldsymbol{b}$ is a two vector.

Proof. The first statement is a straightforward consequence of the group-closure property for affine transformations. Since the two sides Γ and Γ' of the object are related by a reflection (which is affine), and since imaging is assumed to be affine, the two imaged sides γ and γ' are also related by an affine transformation.

In more detail, and to fix notation, object curves Γ and Γ' are related by a reflection, i.e. if $\boldsymbol{X}$ is a point on Γ then there is a point $\boldsymbol{X}'$ on Γ' such that

$$\boldsymbol{X}' = S\boldsymbol{X} + \boldsymbol{B}_{\mathrm{S}}, \tag{A 1}$$

where S is a reflection matrix ($|S| = -1$). Under image transformation (2.4)

$$\boldsymbol{X} = U\boldsymbol{x} + \boldsymbol{B}, \quad \boldsymbol{X}' = U\boldsymbol{x}' + \boldsymbol{B}.$$

Combining these with equation (A 1) gives

$$\boldsymbol{x}' = A\boldsymbol{x} + \boldsymbol{b},$$

where

$$A = U^{-1}SU, \quad \boldsymbol{b} = U^{-1}(S\boldsymbol{B} + \boldsymbol{B}_{\mathrm{S}} - \boldsymbol{B}).$$

■

Property 2. The affine transformation $\{A, \boldsymbol{b}\}$ obeys the following constraints:
(i) $A^2 = I$;
(ii) $|A + I| = 0$; and
(iii) $|A - I| = 0$.

Proof. The first part is straightforward:

$$\begin{aligned} A &= U^{-1}SU \\ A^2 &= U^{-1}S^2U \\ &= U^{-1}U \\ &= I. \end{aligned}$$

The image correspondence $\boldsymbol{x} \leftrightarrow \boldsymbol{x}'$ relates points transformed according to the first part of the theorem. So, $\boldsymbol{x}' = A\boldsymbol{x} + \boldsymbol{b}$ and also $\boldsymbol{x} = A\boldsymbol{x}' + \boldsymbol{b}$. Consequently, applying this transformation to a point $\boldsymbol{x}$ on γ maps it to a point $\boldsymbol{x}'$ on γ', and applying the transformation again maps it back to $\boldsymbol{x}$:

$$\begin{aligned} \boldsymbol{x} &= A(A\boldsymbol{x} + \boldsymbol{b}) + \boldsymbol{b} \\ &= A^2\boldsymbol{x} + (A + I)\boldsymbol{b}. \end{aligned}$$

Since this is true for all $\boldsymbol{x}$ we have the following.
(i) Again $A^2 = I$. From this it follows that A has eigenvalues ± 1.
(ii) $(A + I)\boldsymbol{b} = 0$. Hence, $|A + I| = 0$ since $\boldsymbol{b}$ is non-trivial, and $\boldsymbol{b}$ is an eigenvector of A with eigenvalue -1.
(iii) If $\boldsymbol{x}$ is on the (imaged) symmetry axis then $\boldsymbol{x}' = A\boldsymbol{x} + \boldsymbol{b} = \boldsymbol{x}$ and hence

$$(I - A)\boldsymbol{x} = \boldsymbol{b}. \tag{A 2}$$

This is true for all points on the axis. Consequently, $(A - I)$ must be of rank one (i.e. $|A - I| = 0$) with the kernel of $A - I$ defining the axis direction $\boldsymbol{a}$, i.e. $(A - I)\boldsymbol{a} = 0$, which means that $\boldsymbol{a}$ is an eigenvector of A with eigenvalue 1. ■

Property 3. The matrix A has eigenvectors $\boldsymbol{a}$ and $\boldsymbol{b}$ with eigenvalues $+1$ and -1, respectively. Vector $\boldsymbol{a}$ is parallel to the symmetry axis, vector $\boldsymbol{b}$ is parallel to $\boldsymbol{x}' - \boldsymbol{x}$.

Proof. All that remains to be shown is that $\boldsymbol{b}$ is parallel to $\boldsymbol{x}' - \boldsymbol{x}$. Note, corresponding image points are joined by parallel lines since these lines are images of parallel lines on the object. Recall that the eigenvectors of A are $\boldsymbol{a}$ and $\boldsymbol{b}$ with eigenvalues $+1$ and -1, respectively. $\boldsymbol{a}$ and $\boldsymbol{b}$ span the image, so that for some α, β,

$$\boldsymbol{x}' - \boldsymbol{x} = \alpha\boldsymbol{a} + \beta\boldsymbol{b}.$$

Applying A to both sides, and reversing the order gives

$$\begin{aligned} \alpha\boldsymbol{a} - \beta\boldsymbol{b} &= A\boldsymbol{x}' - A\boldsymbol{x} \\ &= \boldsymbol{x} - \boldsymbol{x}' \\ &= -\alpha\boldsymbol{a} - \beta\boldsymbol{b}, \end{aligned}$$

from which it follows that $\alpha = 0$, so that $\boldsymbol{x} - \boldsymbol{x}'$ is parallel to $\boldsymbol{b}$ as required. ■

Property 4. The transformation has three degrees of freedom. It can be parametrized by a, b_x, b_y, where

$$A = \begin{bmatrix} a & -b_x(1+a)/b_y \\ -b_y(1-a)/b_x & -a \end{bmatrix}, \quad \boldsymbol{b} = \begin{bmatrix} b_x \\ b_y \end{bmatrix}. \tag{A 3}$$

Proof. Solving the equations $|A + I| = 0$ and $|A - I| = 0$ for A gives

$$A = \begin{bmatrix} a & b \\ (1-a^2)/b & -a \end{bmatrix},$$

where a and b are two parameters. This also satisfies $A^2 = I$. The requirement that $A\boldsymbol{b} = -\boldsymbol{b}$ determines b in terms of a, b_x and b_y.

The three parameters represent the symmetry line (two degrees of freedom) and the correspondence direction (one degree of freedom). ■

Property 5. With this notation, the image (skewed) symmetry line is $(1-a)b_y x + (1+a)b_x y - b_x b_y = 0$.

Proof. This is the particular solution of equation (A 2).

Appendix B. Symmetry under projective transformations

In the most general case there is a projective (rather than affine) transformation between object and image planes. As in the affine case, symmetry in the object plane constrains the transformation between image curves.

Theorem B.1. *Suppose two curves γ and γ', as in figure 2, are the images of two corresponding sides of a planar object with bilateral symmetry. Suppose further that image projection can be represented by a projective transformation. Then the transformation between γ and γ' has the following properties.*

Property 1. γ and γ' are related by a projective transformation. That is, if $\boldsymbol{x}$ is a point on γ then there is a point $\boldsymbol{x}'$ on γ' such that

$$\boldsymbol{x}' = T\boldsymbol{x}, \tag{B 1}$$

where T is a non-singular 3×3 matrix, and $\boldsymbol{x}$ and $\boldsymbol{x}'$ are homogeneous three vectors.

Property 2. The projective transformation T satisfies the following constraints:

(i) $T^2 = kI$, where k is a scalar; and

(ii) the fixed points of T are: a line of fixed points; and a fixed point not on the line (through which there is a pencil of fixed lines).

A projection with these properties is a collineation of period two, also known as a two cyclic homography, and a planar harmonic homology (Springer 1964).

Property 3. The matrix T has eigenvectors $\{\boldsymbol{e}_1, \boldsymbol{e}_2, \boldsymbol{e}_3\}$. Two of the eigenvalues, corresponding to $\boldsymbol{e}_2$ and $\boldsymbol{e}_3$, say, are equal. The third, corresponding to $\boldsymbol{e}_1$ is distinct and non-zero. The symmetry axis is given by the line $\boldsymbol{e}_2 \times \boldsymbol{e}_3$. Corresponding points, $\boldsymbol{x}'$ and $\boldsymbol{x}$, are collinear with $\boldsymbol{e}_1$. The line $\boldsymbol{x}'$, $\boldsymbol{x}$ intersects the symmetry axis at a point $\boldsymbol{x}_I$ say, and the four collinear points $\boldsymbol{x}$, $\boldsymbol{x}_I$ $\boldsymbol{x}'$ and $\boldsymbol{e}_1$ have a harmonic cross-ratio.

Property 4. The transformation has four degrees of freedom. It can be determined from two correspondences.

Proof. Property 1. This follows from group closure under projective transformations. If the object reflection is the projective transformation $\boldsymbol{X}' = S\boldsymbol{X}$, and the object to image projection is

$$\boldsymbol{x} = U\boldsymbol{X}, \quad \boldsymbol{x}' = U\boldsymbol{X}',$$

then the transformation between γ and γ' is given by the conjugate projectivity:

$$T = USU^{-1}. \qquad \blacksquare$$

Property 2.(i). We have the image correspondence $\boldsymbol{x} \leftrightarrow \boldsymbol{x}'$, so that $\boldsymbol{x}' = T\boldsymbol{x}$ and also $\boldsymbol{x} = T\boldsymbol{x}'$. Consequently, applying this transformation to a point $\boldsymbol{x}$ on γ maps it to a point $\boldsymbol{x}'$ on γ', and applying the transformation again maps it back to $\boldsymbol{x}$:

$$\boldsymbol{x} = T(T\boldsymbol{x}) = T^2\boldsymbol{x}$$

Since this is true for all $\boldsymbol{x}$, $T^2 = kI$.

(ii). This can be proved analytically but is seen most simply by considering the projection geometry in figure 26. On the object, points on the symmetry axis are mapped by the reflection to themselves – so this is a line of fixed points. Correspondingly, the imaged symmetry line is a line of fixed points. The only other fixed point (not on this line) is the point at infinity (where parallel lines joining corresponding points intersect). In the image, this is transformed to the vanishing point of the lines joining corresponding points. $\blacksquare$

Property 3. For a line $\boldsymbol{l}$ of fixed points, two degenerate eigenvectors must lie on the line. To see this, represent a point on the line as $\boldsymbol{x} = \mu\boldsymbol{e}_2 + \nu\boldsymbol{e}_3$. Then

$$\boldsymbol{x}' = T(\mu\boldsymbol{e}_2 + \nu\boldsymbol{e}_3) = \lambda(\mu\boldsymbol{e}_2 + \nu\boldsymbol{e}_3),$$

(where $\boldsymbol{e}_2$ and $\boldsymbol{e}_3$ are eigenvectors of T with eigenvalue λ) which is the same point.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

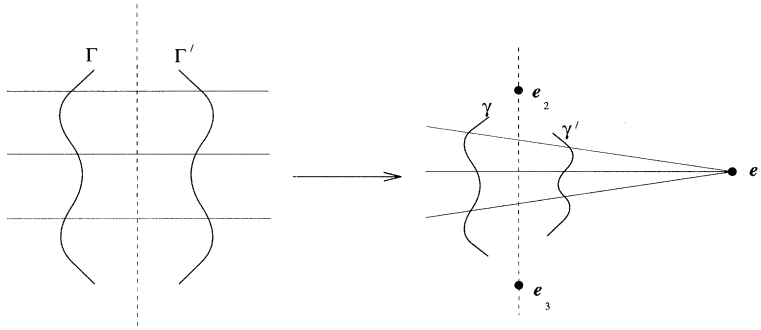

Figure 26. Under a projective transformations parallel object correspondences converge to a vanishing point.

Thus, the symmetry line is given by $\boldsymbol{e}_2 \times \boldsymbol{e}_3$. The vanishing point corresponds to $\boldsymbol{e}_1$, the other fixed point. A similar argument to the above shows that any line passing through this point is a fixed line under T. On the object, corresponding points are collinear with this point at infinity (since the lines joining corresponding points are all parallel) so in the image, $\boldsymbol{x}'$ and $\boldsymbol{x}$ are collinear with $\boldsymbol{e}_1$.

On the object, the four points, $\boldsymbol{X}'$, $\boldsymbol{X}$, the intersection of their common line with the symmetry axis and the point at infinity, have a harmonic cross-ratio. This is preserved by projectivities. Hence, the four image points, $\boldsymbol{x}'$, $\boldsymbol{x}$, $\boldsymbol{e}_1$ and the intersection of their common line with the symmetry axis, have a harmonic cross-ratio. ■

Property 4. The transformation has four degrees of freedom. These correspond to the symmetry axis (two degrees of freedom) and the vanishing point (two degrees of freedom). Two point correspondences, $\boldsymbol{x}_1 \leftrightarrow \boldsymbol{x}_1'$ and $\boldsymbol{x}_2 \leftrightarrow \boldsymbol{x}_2'$, give four constraints:

$$\boldsymbol{x}_1' = T\boldsymbol{x}_1, \quad \boldsymbol{x}_2' = T\boldsymbol{x}_2, \quad \boldsymbol{x}_1 = T\boldsymbol{x}_1', \quad \boldsymbol{x}_2 = T\boldsymbol{x}_2',$$

which is sufficient to determine T. ■

Projected symmetry axis. The above results provide an 'optimal' algorithm for extracting the projected symmetry axis – namely determine a T which best maps one side of the contour to the other, and compute the axis from the eigenvectors. An alternative simple method arises by noting that any projectively covariant construction on the object which generates points on the symmetry axis, can be used to determine points on the symmetry axis in the image. For example, find corresponding pairs of distinguished points on each side of the outline, say a corresponding to a′, b corresponding to b′. Then line pairs {ab′, a′b} and {ab, a′b′} both intersect on the symmetry axis. These two intersections determine the line. In the absence of measurement noise both the 'optimal' and simple method will produce the same line.

References

Beiderman, I. 1987 Recognition-by-components: a theory of human image understanding. *Psychol. Rev.* **94**, 115–147.

Blake, A. & Marinos, C. 1990 Shape from texture: estimation, isotropy and moments. *Artif. Intell.* **45**, 323–380.

Blake, A. & Taylor, M. 1993 Grasping visual symmetry. *Proc ICCV4*, 724–733.

Blum, H. 1973 Biological shape and visual science (part I). *J. theor. Biol.* **38**, 205–287.

Blum, H. & Nagel, R. N. 1978 Shape description using weighted symmetric axis features. *Pattern Recognition* **10**, 167–180.

Brady, J. M. & Asada, H. 1984 Smoothed local symmetries and their implementation. *Int. J. Robotics Res.* **3**, 36–61.

Brady, J. M. & Yuille, A. 1984 An extremum principle for shape from contour. *PAMI* **6**, 288–301.

Canny, J. F. 1986 Finding edges and lines in images. *PAMI* **8**, 6.

Connell, J. H. & Brady, J. M. 1987 Generating and generalising models of visual objects. *Artif. Intell.* **31**, 159–184.

Corballis, M. C. 1988 Recognition of disoriented shapes. *Psychol. Rev.* **95**, 115–123.

Fawcett, R., Zisserman, A. & Brady, J. M. 1993 Extracting structure from single affine views of 3D point sets with one or two bilateral symmetries. *Proc. BMVC.* (To appear.)

Field, M. & Golubitsky, M. 1992 *Symmetry in chaos: a search for pattern in mathematics, art and nature.* Oxford University Press.

Fleck, M. 1985 Local rotational symmetries. *MIT Rep.* AI-TR-852.

Giblin, P. J. & Brassett S. A. 1985 Local symmetry of plane curves. *Am. math. Monthly* **92**, 689–707.

Grimson, W. E. L. 1990 *Object recognition by computer, the role of geometric constraints.* MIT Press.

Horaud, R. & Brady, J. M. 1988 On the geometric interpretation of image contours. *Artif. Intell.* **37**, 333–353.

Kanade, T. 1981 Recovery of three dimensional shape of an object from a single view. *Artif. Intell.* **17**, 409–460.

Lamdan, Y., Schwartz, J. & Wolfson, H. 1988 Object recognition by affine invariant matching. *Proc. CVPR*, 335–344.

Lowe, D. G. 1985 *Perceptual organization and visual recognition.* Deventer: Kluwer.

Mundy, J. & Zisserman, A. (eds) 1992 *Geometric invariance in computer vision.* MIT Press.

Reid, I. 1991 Recognising parameterized models from range data. D.Phil. thesis, Department of Engineering Science, University of Oxford.

Rom, H. & Medioni, G. 1993 Hierarchical decomposition and axial representation of shape. *PAMI* **15**, 973–981.

Rothwell, C. A., Zisserman, A., Forsyth, D. A. & Mundy, J. L. 1992*a* Canonical frames for planar object recognition. In *Proc. ECCV2*, pp. 757–772. Berlin: Springer.

Rothwell, C. A., Zisserman, A., Marinos, C., Forsyth, D. A. & Mundy, J. L. 1992*b* Relative motion and pose from arbitrary plane curves. *Image and Vision Computing* **10**, 250–262.

Rothwell, C. A., Zisserman, A., Mundy, J. L. & Forsyth, D. A. 1992*c* Efficient model library access by projectively invariant indexing. In *Proc. CVPR*, pp. 109–114.

Scott, G. L., Turner, S. & Zisserman, A. 1989 Using a mixed wave/diffusion process to elicit the symmetry set. *Image and Vision Computing* **7**, 63–70.

Springer, C. E. 1964 *Geometry and analysis of projective spaces.* San Francisco, CA: Freeman.

Strachan, N. J. 1993 Recognition of fish species by colour and shape. *Image and Vision Computing* **11**, 2–10.

Tsai, R. Y. 1986 An efficient and accurate camera calibration technique for 3D machine vision. *Proc. IEEE CVPR 86.*

Van Gool, L. J., Moons, T., Pauwels, E. & Oosterlinck, A. 1992 Semi differential invariants. In Mundy, J. & Zisserman, A. 1992 *Geometric invariance in computer vision.* MIT Press.

Wagemans, J. 1993 Skewed symmetry: A nonaccidental property used to perceive visual forms. *J. Exp. Psychol.: Human Perception and Performance* **19**, 1–17.

Weyl, H. 1952 *Symmetry.* Princeton University Press.

Received 25 June 1993; accepted 4 July 1994

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

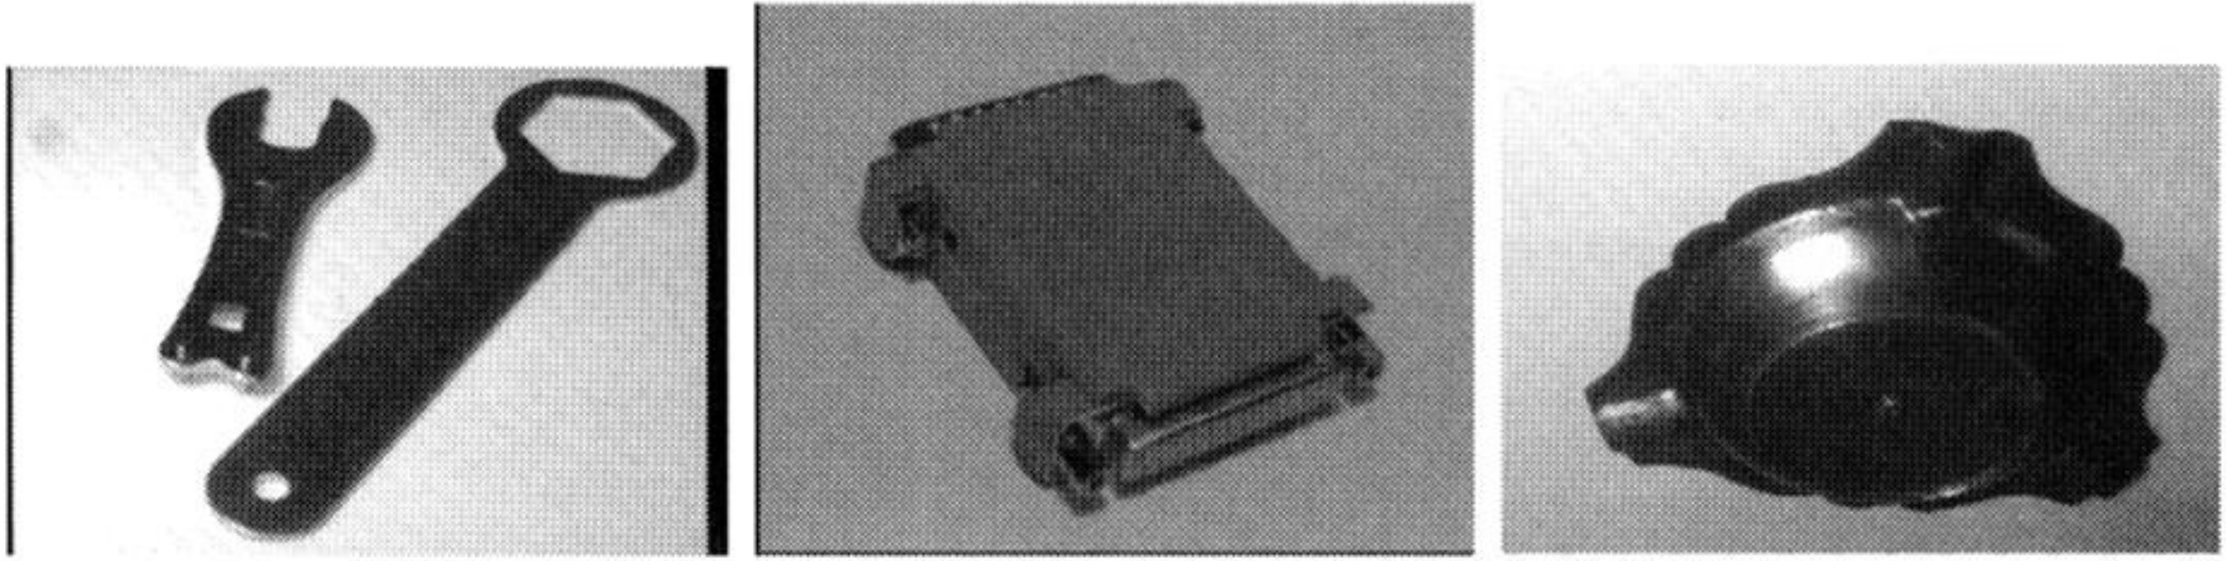

igure 1. Examples of skewed symmetry. The objects differ in the number and disposition of the ymmetry axes: (*a*) each object has a single bilateral symmetry; (*b*) the object has two orthogonal ilateral symmetries; and (*c*) the object has three bilateral symmetries, but the symmetry axes re not orthogonal. In each case the symmetry is global, and *on the object* lines between points elated by the reflectional symmetry are orthogonal to the axis of reflection. However, in the nage these lines and the axis are not orthogonal, in general, but skewed.

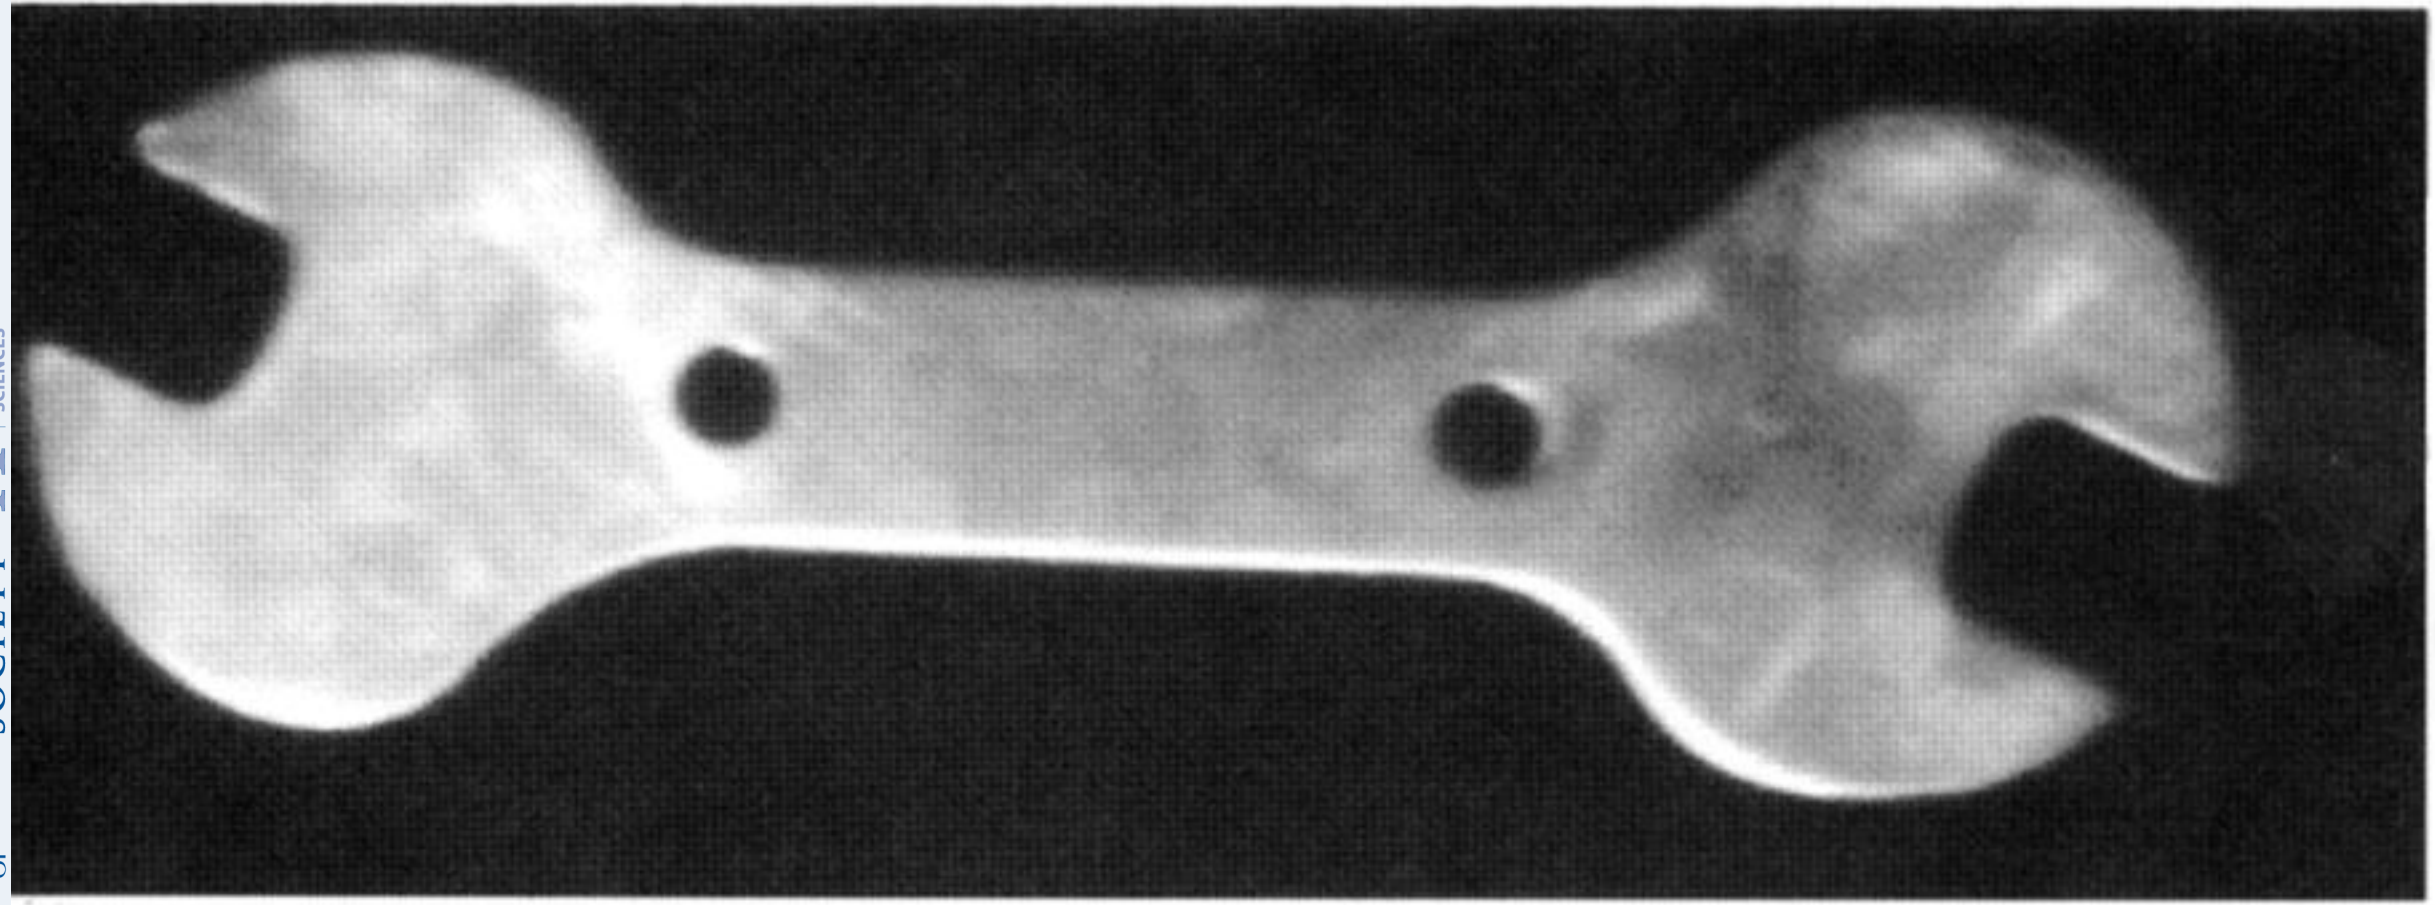

Figure 4. Affine view of spanner.

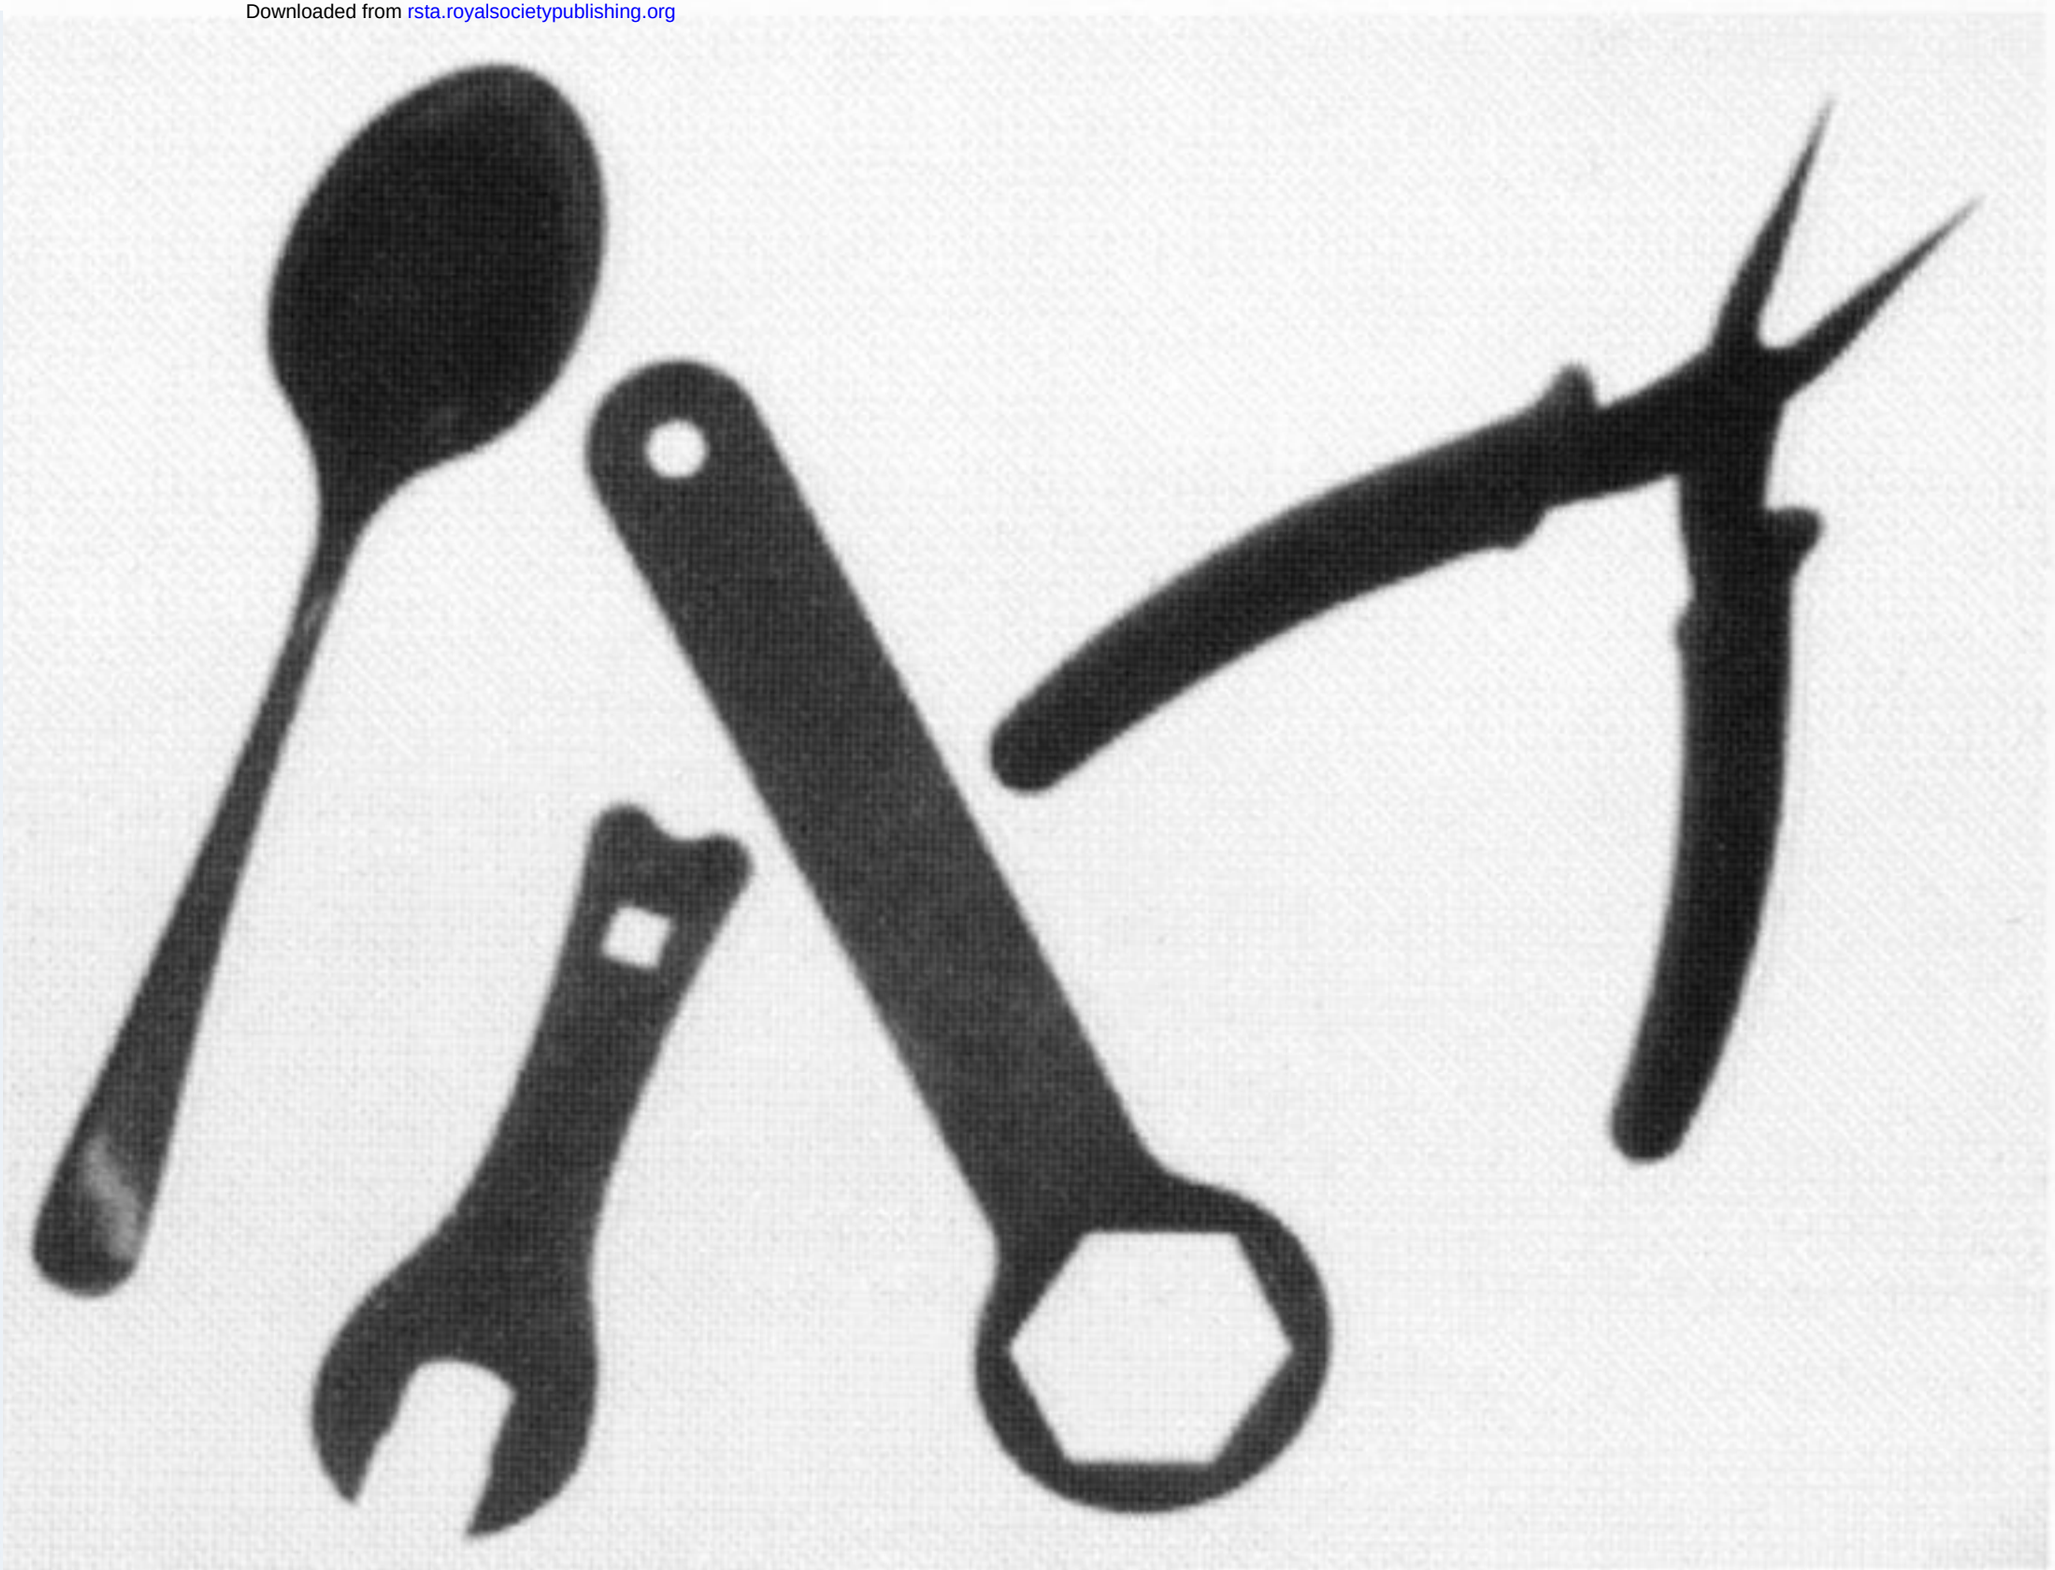

Figure 6. Image of four objects with bilateral symmetries.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

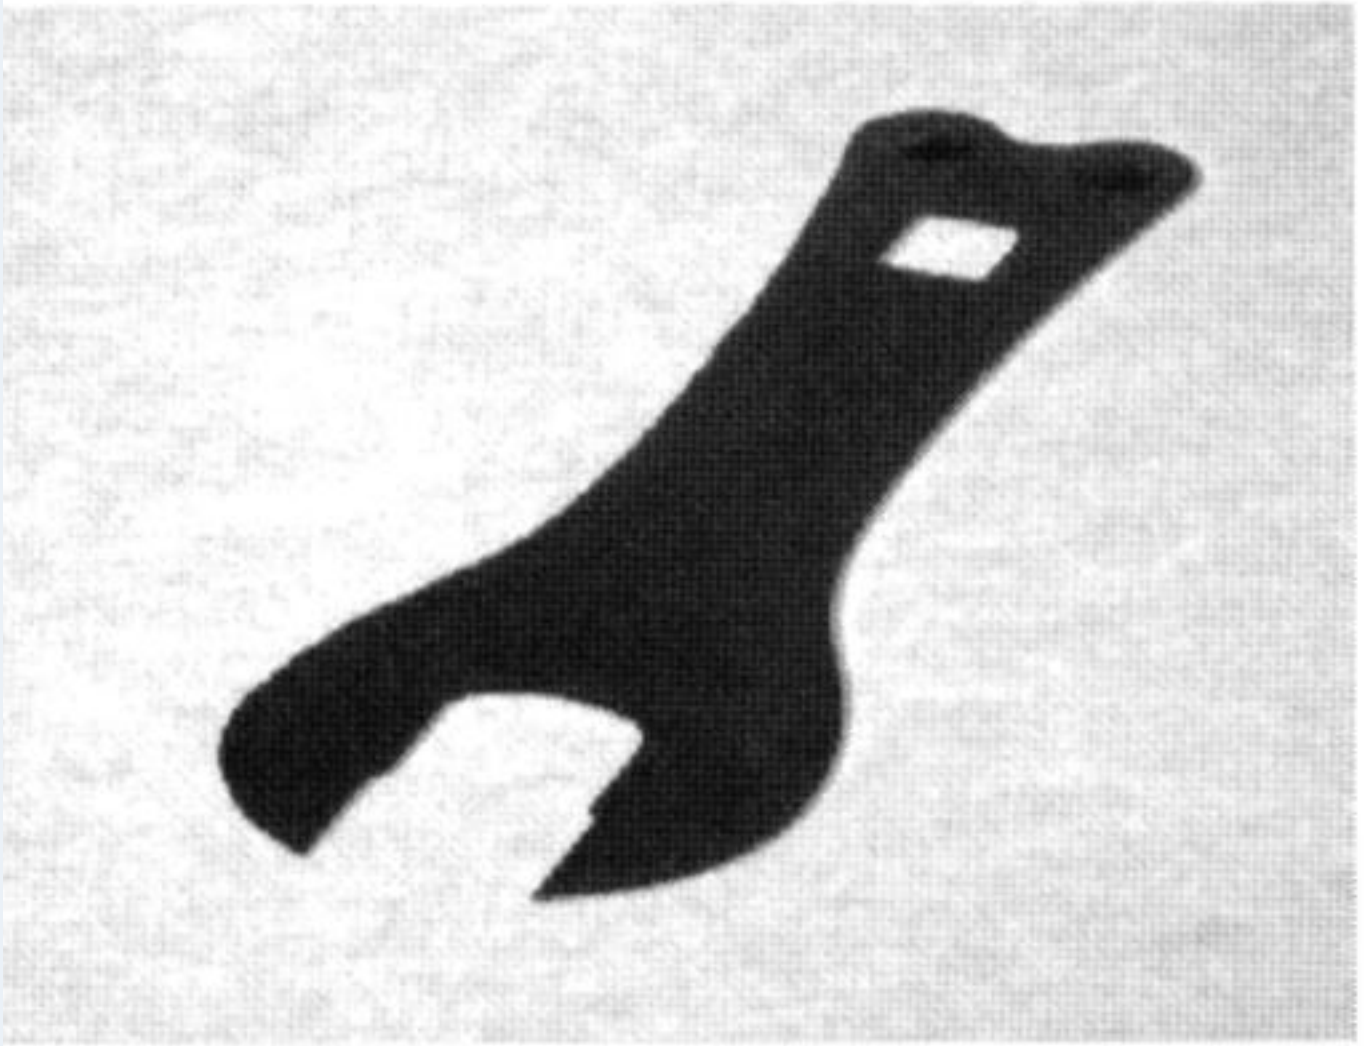

Figure 8. Affine view of symmetric spanner.

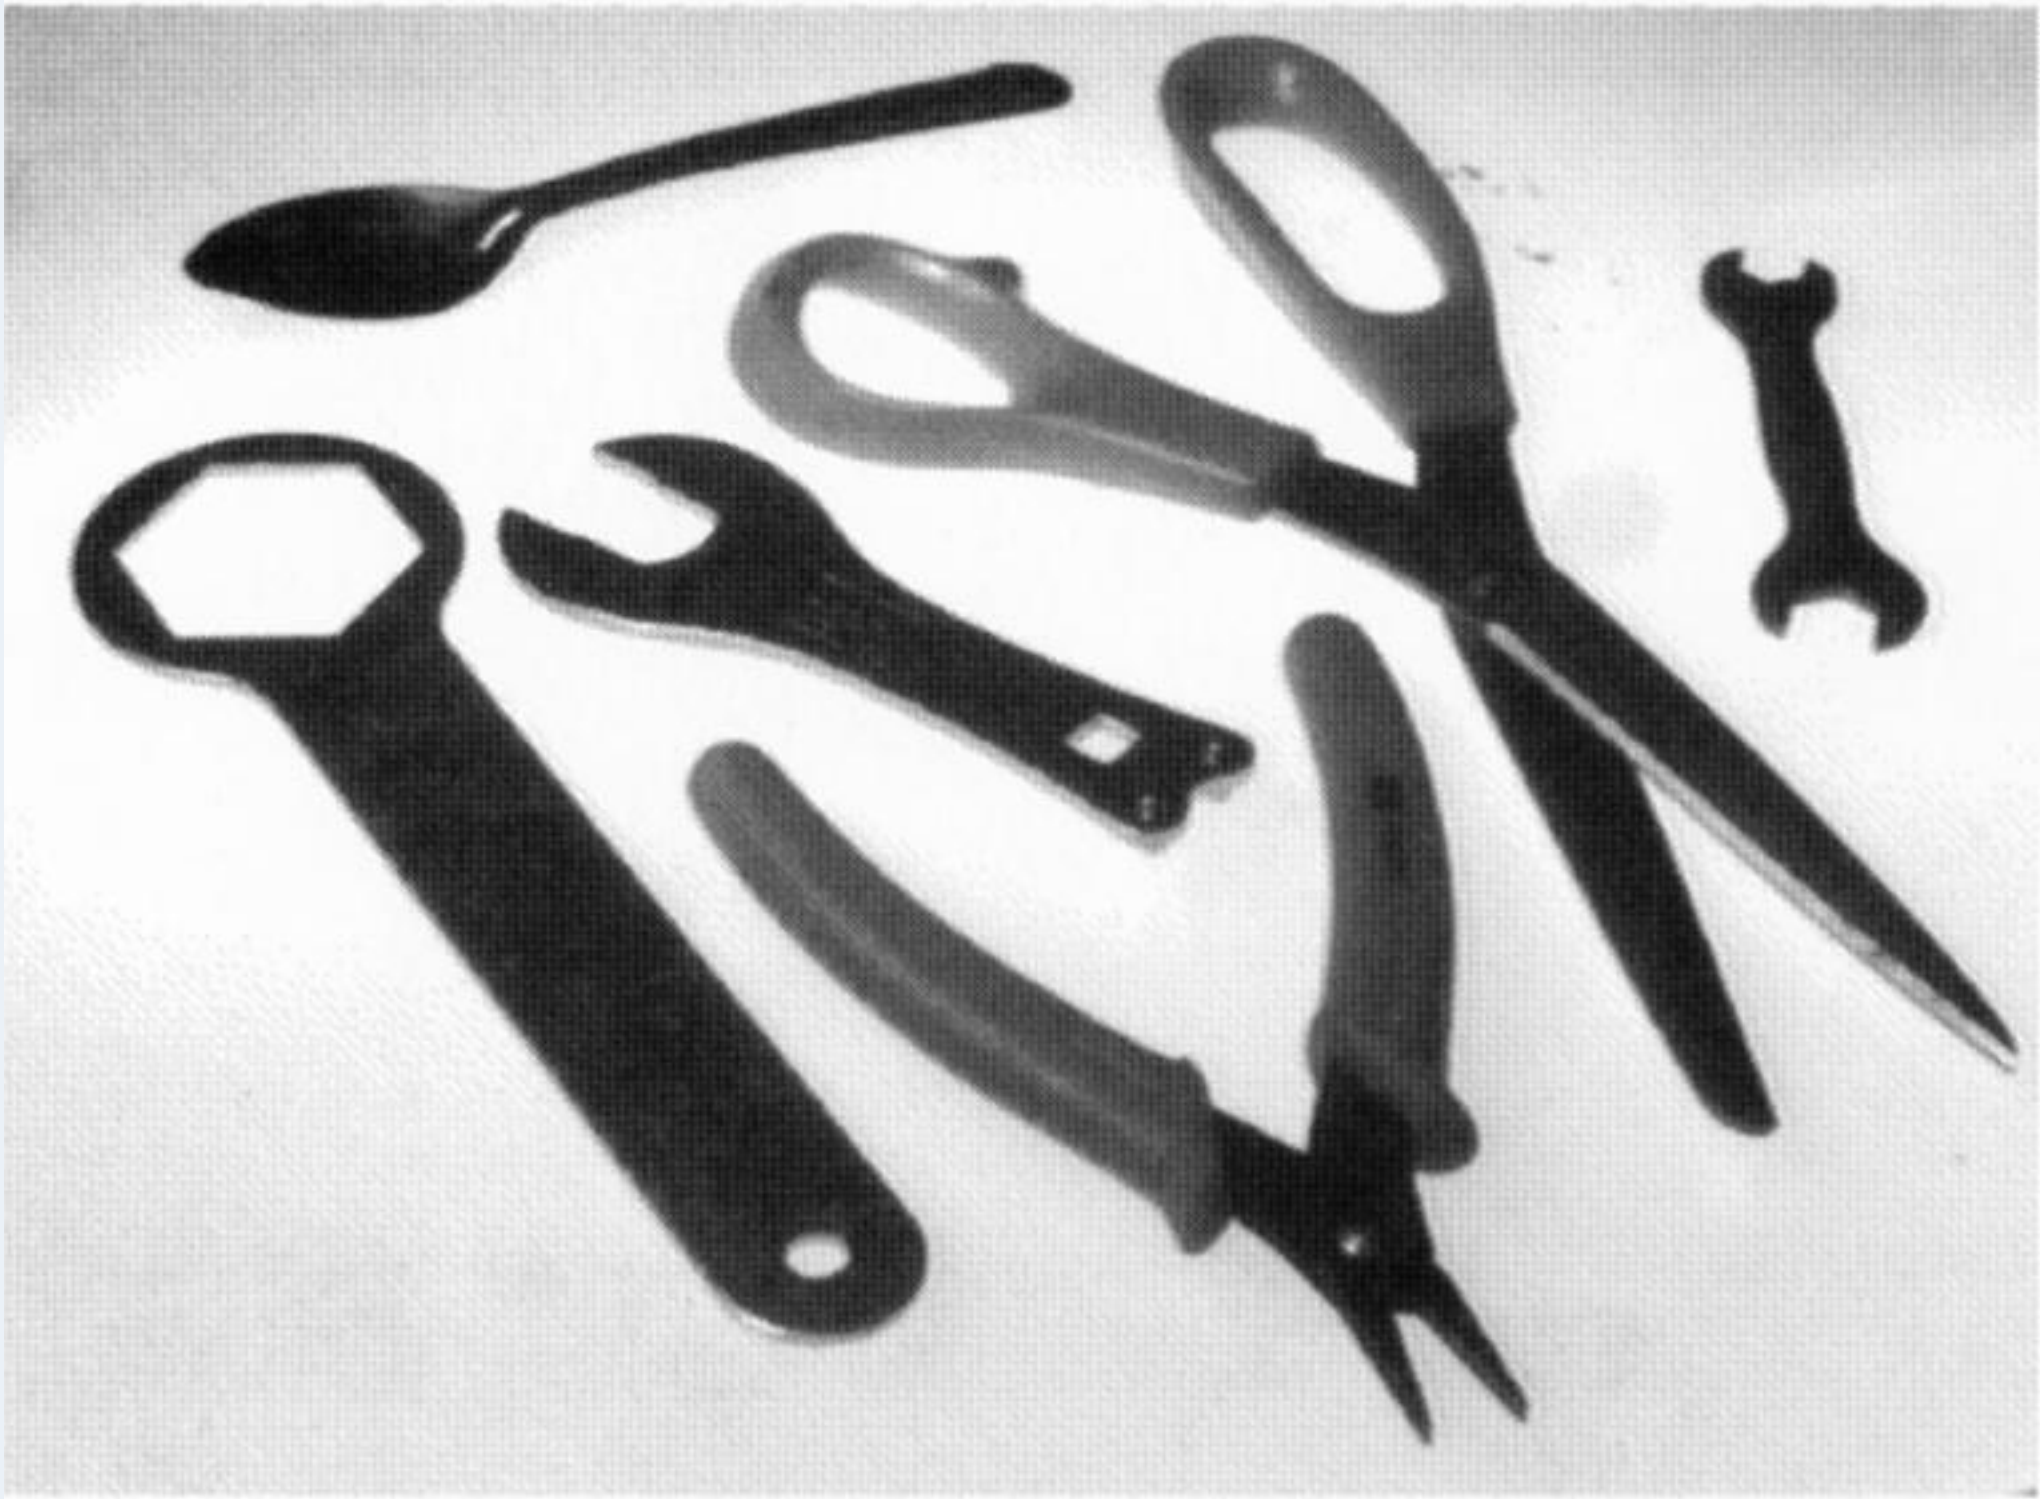

igure 11. Affine scene of globally, partially and non-symmetric objects.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

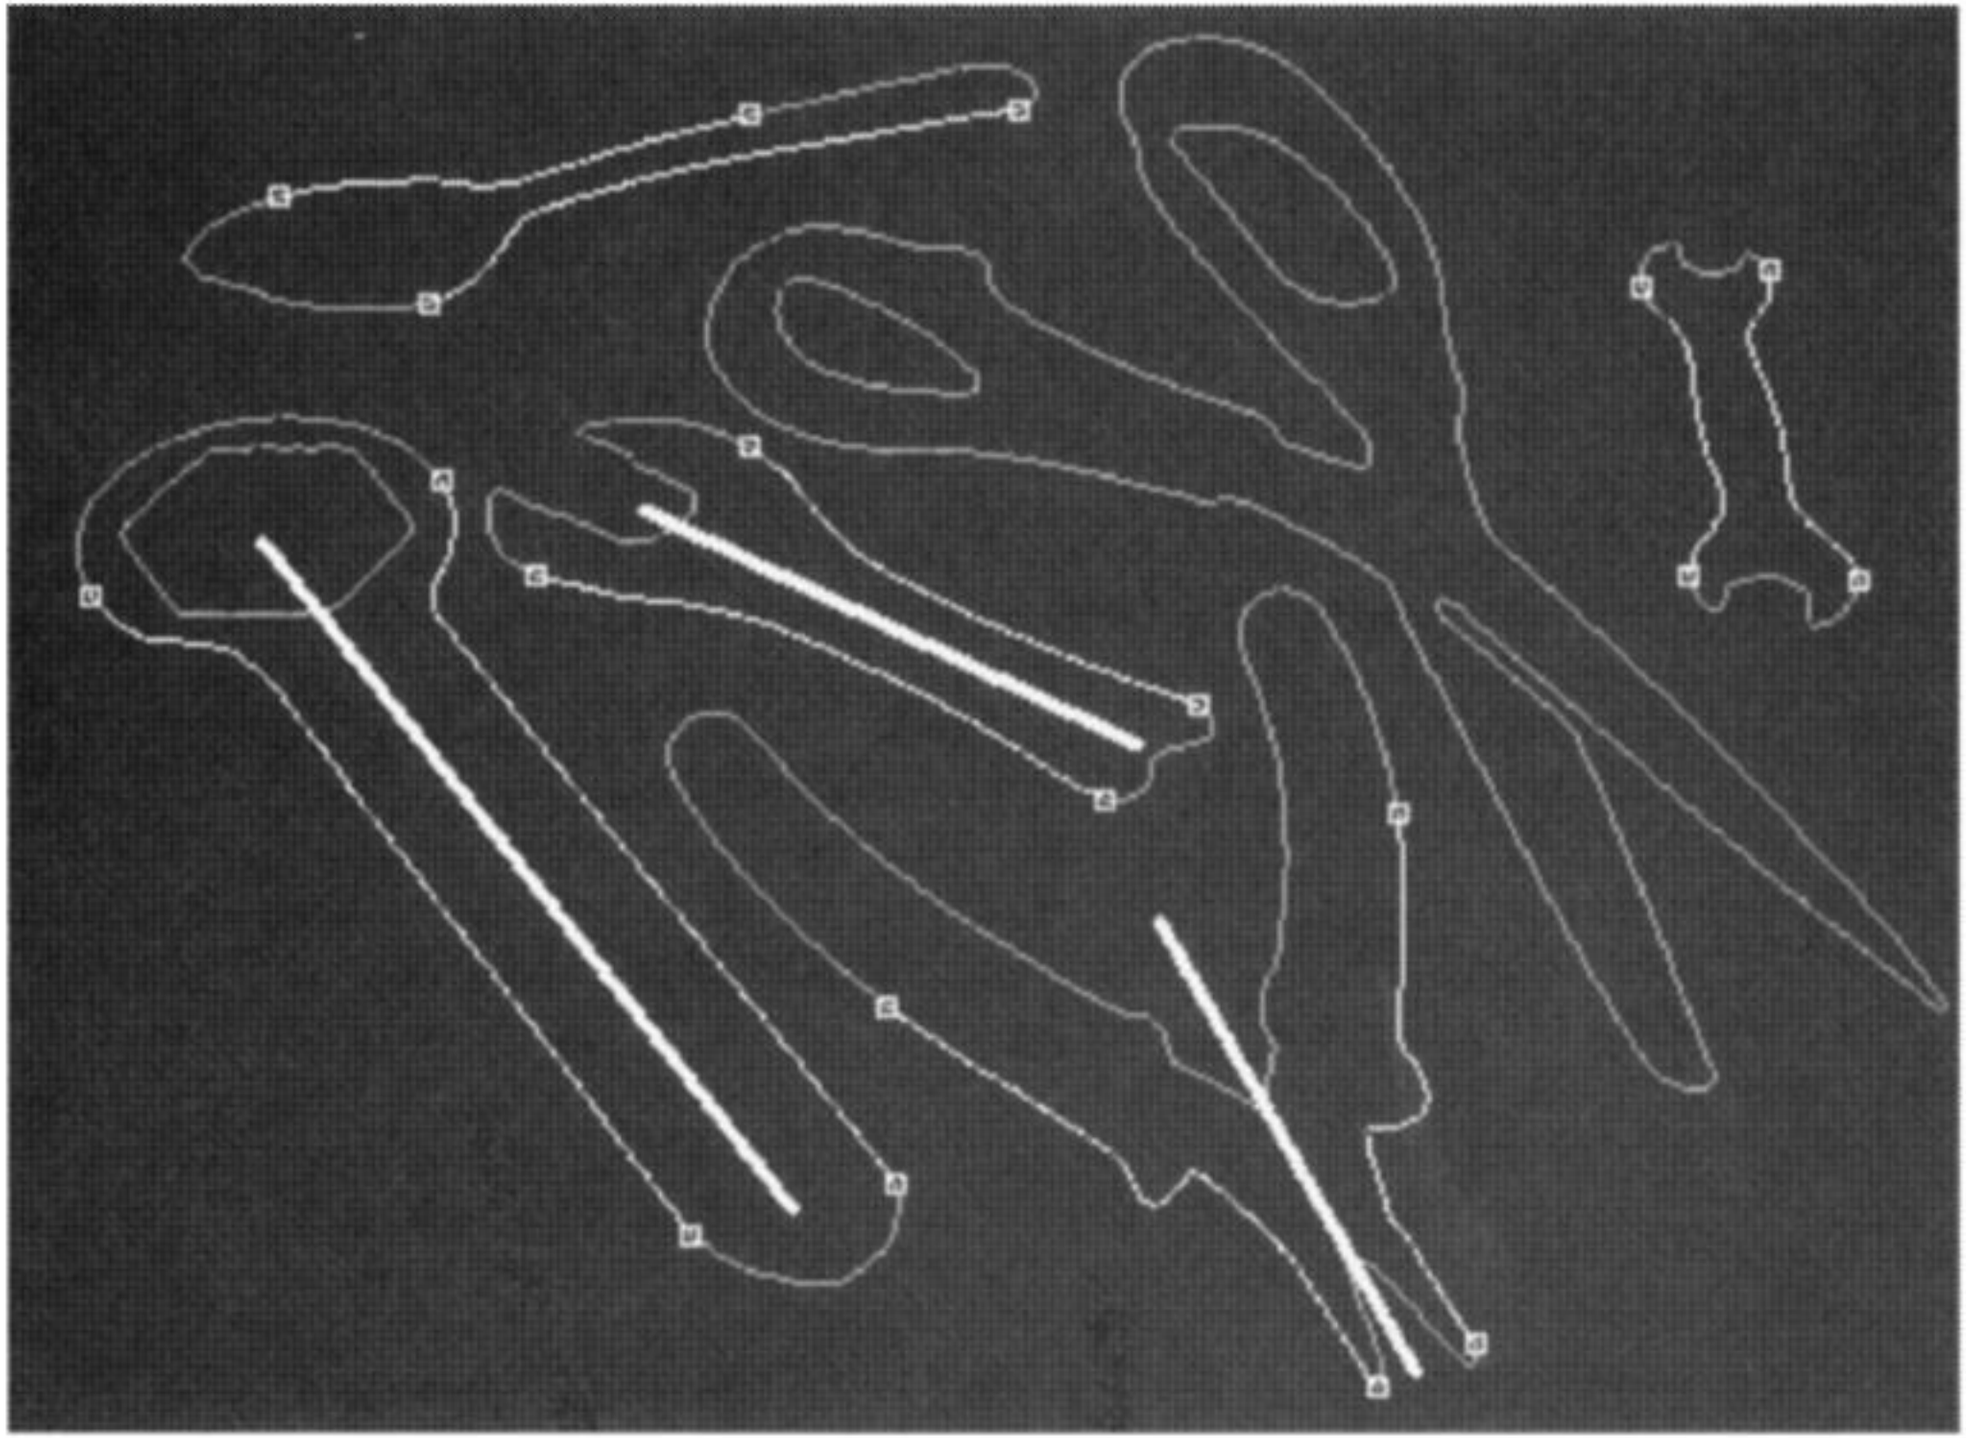

Figure 12. Outline curves, concavity entrance and exit point for the objects in figure 11. Symmetry lines are only drawn where the object is determined to be globally symmetric.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

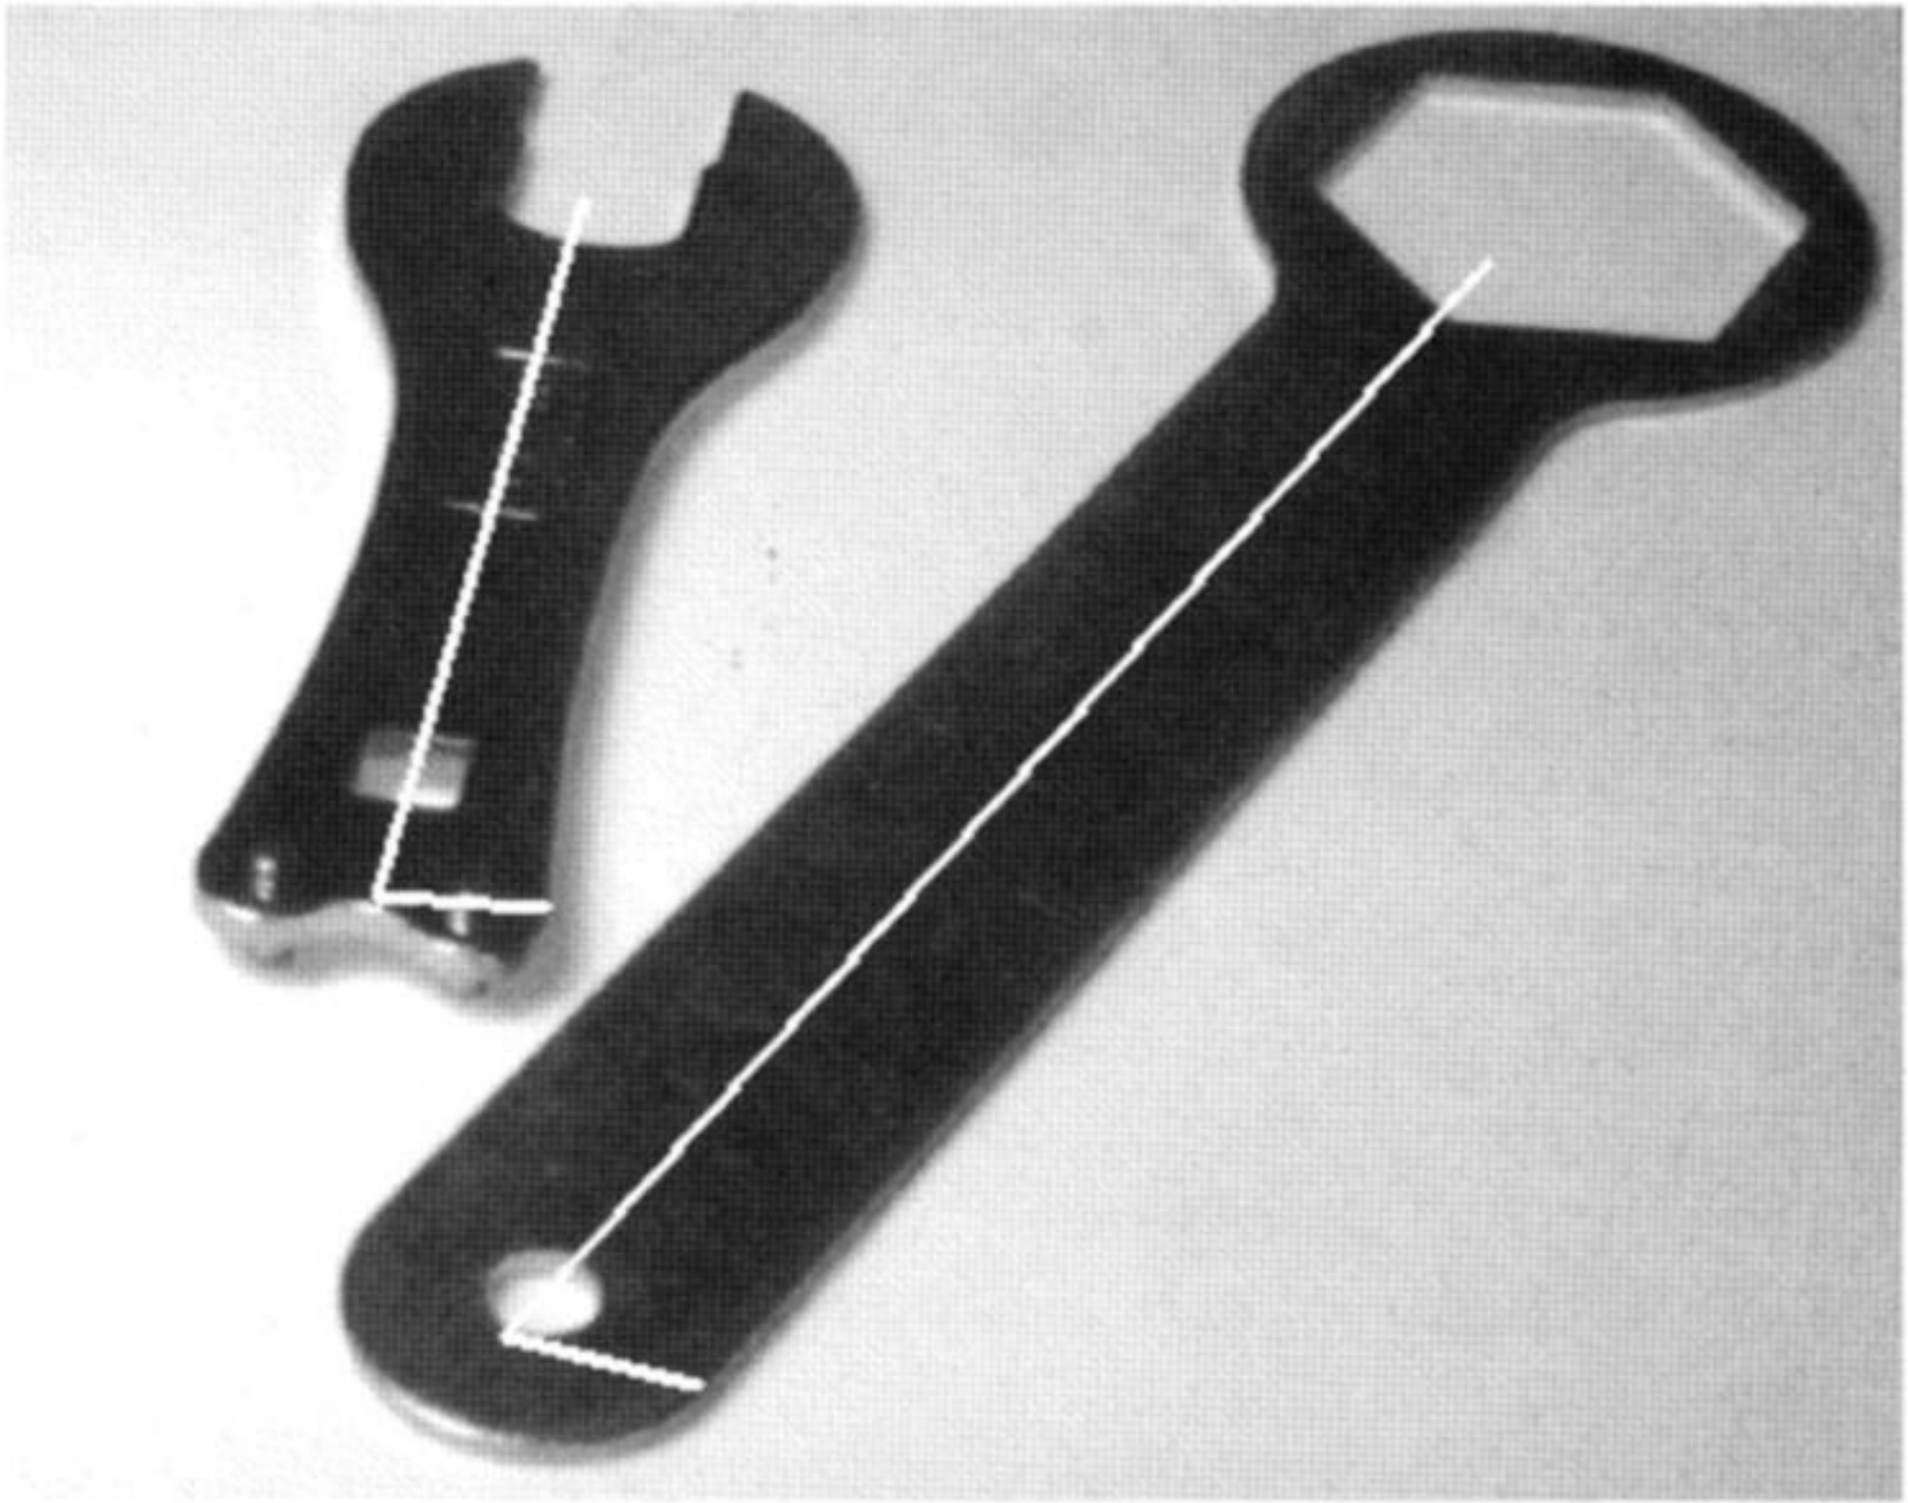

Figure 14. Affine view of two objects with superimposed guiding vectors. The vectors are clearly not perpendicular.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

igure 16. Affine image of an object with multiple local symmetries.

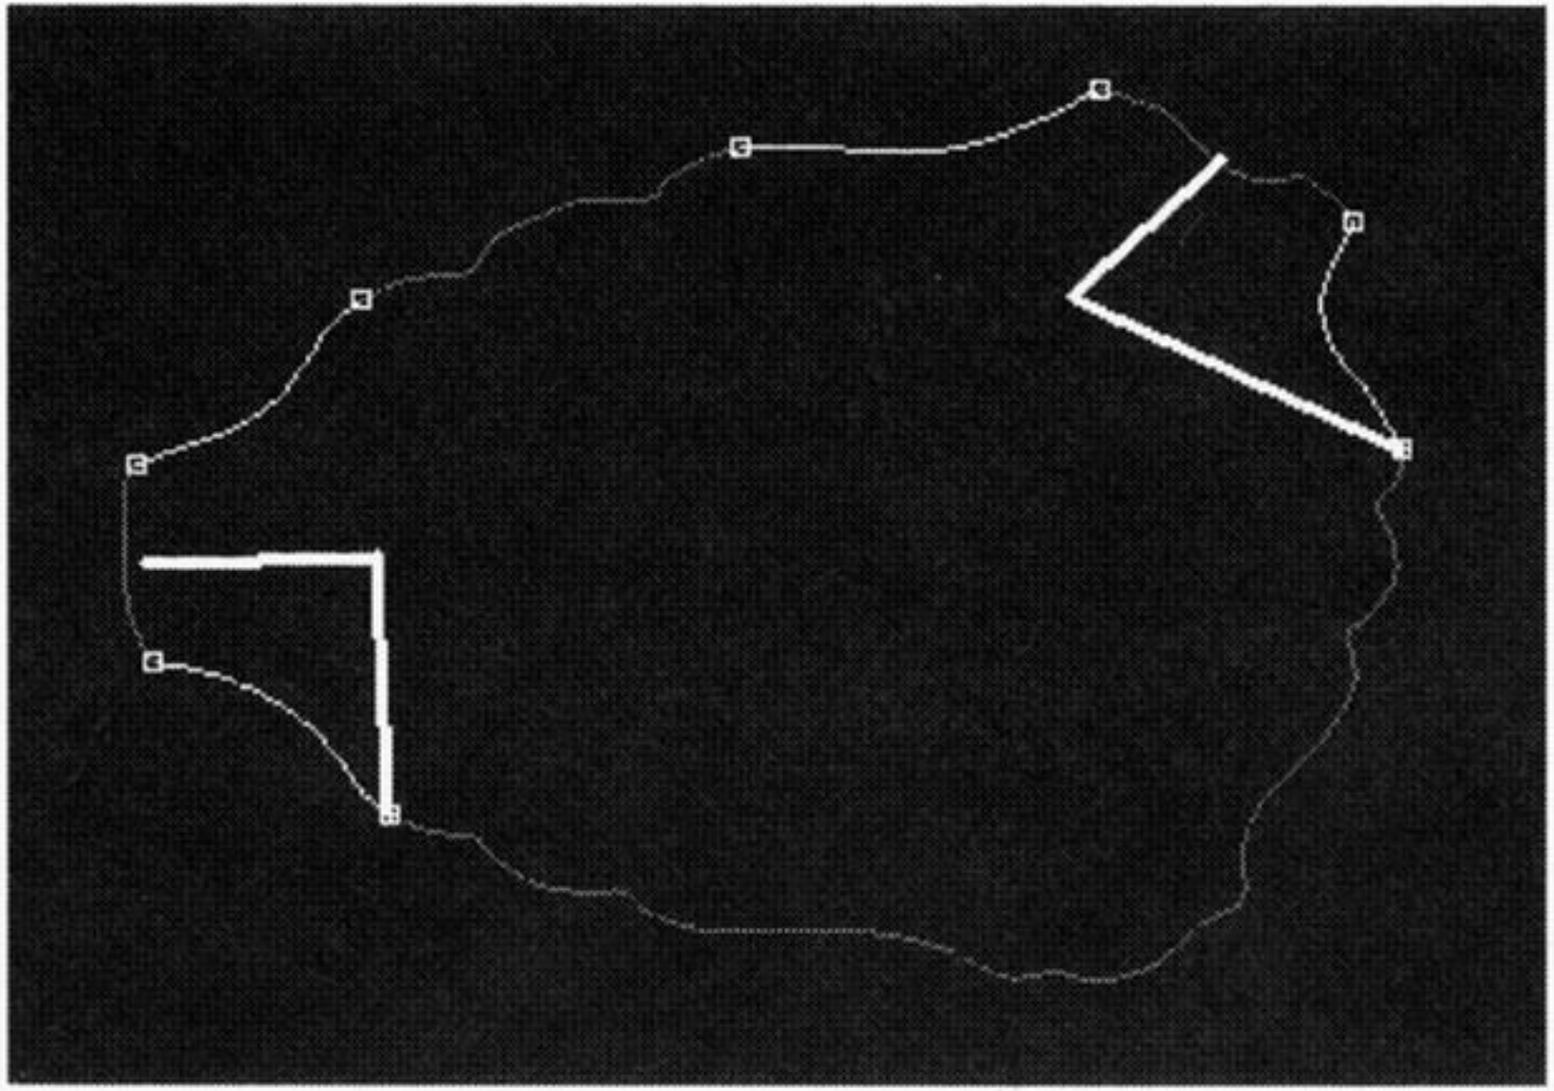

igure 17. Local symmetries and matching concavities extracted from figure 16. Note that local symmetries considered are neither parallel nor perpendicular.

Figure 19. Affine scene of multiple objects.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

igure 21. Skewed image of hex spanner with superimposed guiding vectors.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

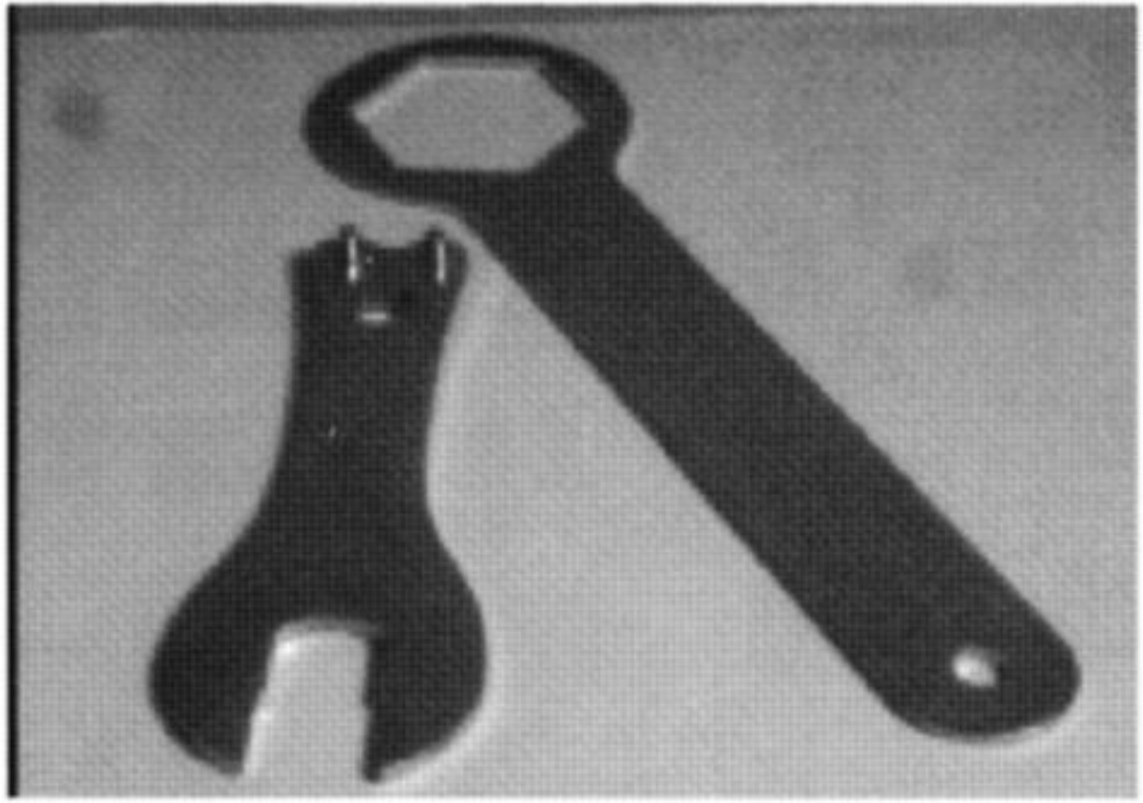

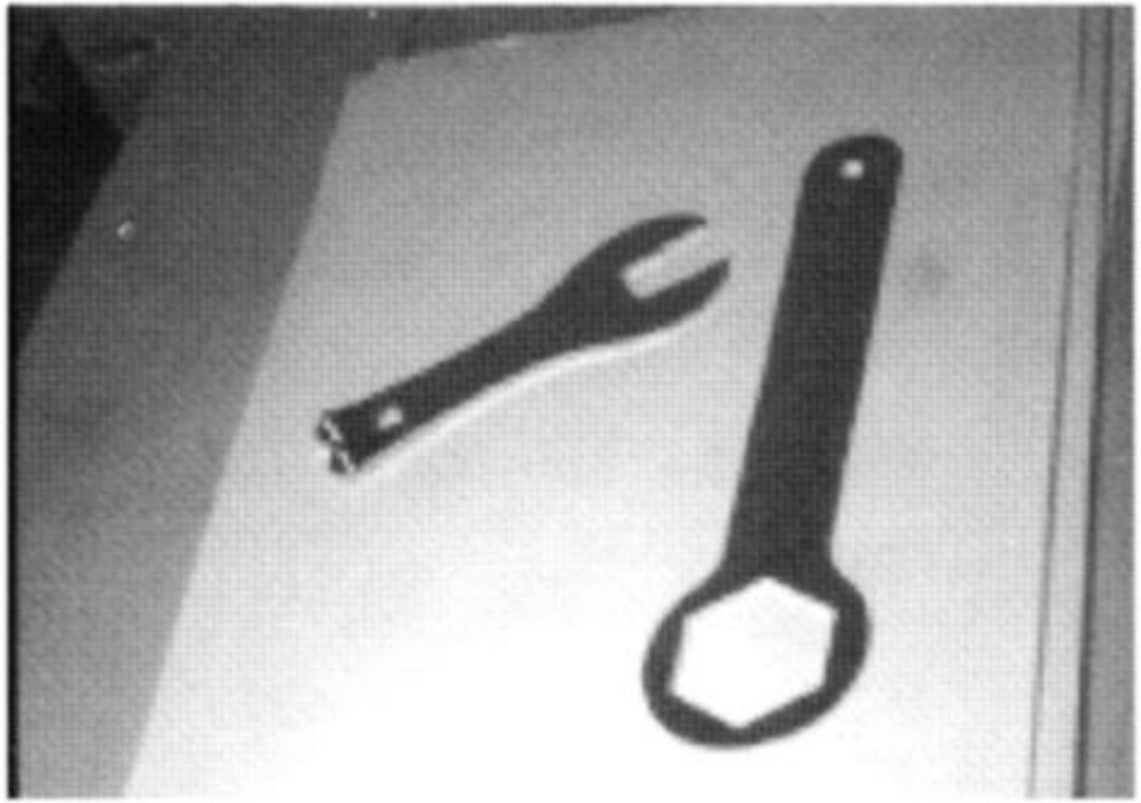

igure 24. Examples of images used to calculate σ and τ for each of the camera orientations. or each orientation slant and tilt is recovered for three different object arrangements. Results e given in table 6.

Figure 25. Affine image of two non-coplanar objects.