Geometric Saliency of Curve Correspondences and Grouping of Symmetric Contours

Tat-Jen Cham and Roberto Cipolla

Department of Engineering
University of Cambridge
Cambridge CB2 1PZ, England

Abstract. Dependence on landmark points or high-order derivatives when establishing correspondences between geometrical image curves under various subclasses of projective transformation remains a shortcoming of present methods. In the proposed framework, geometric transformations are treated as smooth functions involving the parameters of the curves on which the transformation basis points lie. By allowing the basis points to vary along the curves, hypothesised correspondences are *freed* from the restriction to fixed point sets. An optimisation approach to localising *neighbourhood-validated* transformation bases is described which uses the deviation between projected and actual curve neighbourhood to iteratively improve correspondence estimates along the curves. However as transformation bases are inherently localisable to different degrees, the concept of *geometric saliency* is proposed in order to quantise this localisability. This measures the sensitivity of the deviation between projected and actual curve neighbourhood to perturbation of the basis points along the curves. Statistical analysis is applied to cope with image noise, and leads to the formulation of a *normalised basis likelihood*. Geometrically salient, neighbourhood-validated transformation bases represent hypotheses for the transformations relating image curves, and are further refined through curve support recovery and geometrically-coupled active contours. In the thorough application of this theory to the problem of detecting and grouping *affine* symmetric contours, good preliminary results are obtained which demonstrate the independence of this approach to landmark points.

1 Introduction

Establishing correspondences between points in images or 3D structure models in databases is often an integral part of solving key problems in computer vision. The choice of correspondences is largely determined by the localisability of the points, and methods such as cornerness detectors [7] have been proposed to extract these interest points from models and images. In cases such as stereo vision and motion, correspondences between interest points are hypothesised through the use of intensity correlation techniques [1] with the assumption that point features in the different images do not differ much.

Methods based solely on the geometry of the contours are certainly more applicable in problems such as object recognition and pose identification. Various strategies for identifying salient points and regions on contours have been suggested [10, 5] but they are seldom invariant beyond similarity transformations. More appropriately, *landmark* points such as corners and inflection points are used, eg. to compute algebraic invariants for comparison purposes [15]. Although less computationally expensive than Hough methods [1] and more stable than using high-ordered differential invariants [8], the shortfall of this approach lies in the dependence on the availability and accurate localisation of landmark points.

Similar difficulties are also encountered in the detection of image contours which are geometrically related by properties such as parallel or skewed symmetry [11, 9, 6]. The recovery of the geometric transformations in these situations will facilitate the grouping of geometrically related contours which may be fragmented due to occlusion and background clutter. For general classes of objects, the grouping provided by geometry is an enhancement over traditional methods of perceptual grouping. Results by West and Rosin [13], Zerroug and Nevatia [14] and Zisserman *et al.* [16] for the recovery of the contours of surfaces of revolution and canal surfaces are promising, although the correspondences are found only if specific image features are located.

In this paper, a general theory of geometric curve-based correspondence is proposed. We initially discuss general approaches to establishing curve correspondences, followed by introducing transformation bases as functions of curve parameters. We further outline a process to localise transformation bases in order to satisfy the curve neighbourhood at the basis points. The concept of *geometric saliency* is proposed to evaluate the discriminatory power of different transformation bases, and statistical analysis is carried out to identify salient bases in the presence of noise, resulting in the derivation of a *normalised basis likelihood*. Finally, the theory is applied to the problem of detecting and grouping affine symmetric curves.

2 Theory

2.1 Formation and Parameterisation of Transformation Bases

The unique point-to-point correspondence between two sets of image curves related by up to a *2D projective* transformation may be obtained once an initial N point correspondences in general position have been established together with P non-invariant contour derivatives at these points, where $2N + P \geq d$ and $d \leq 8$ is the number of degrees of freedom available in the transformation class. The collections of minimum number of pairs of corresponding points and associated derivatives required to fix the free parameters (ie. the equality $2N + P = d$ holds) form the **bases** of the transformation (Fig. 1). See [12] for an in-depth discussion.

Associated with every transformation basis are matrices which relate corresponding points in the basis. Since in our framework all points lie on curves, we designate these pairs of corresponding points as $(\boldsymbol{x}_{1j}(t_{1j}), \boldsymbol{x}_{2j}(t_{2j})), j = 1, .., N$, where t_{ij} is the parameterisation for the curve on which $\boldsymbol{x}_{ij}$ lie. Then

$$\boldsymbol{A}_{12} : \boldsymbol{x}_{1j}(t_{1j}) \longmapsto \boldsymbol{x}_{2j}(t_{2j}), \qquad j = 1, .., N \tag{1}$$

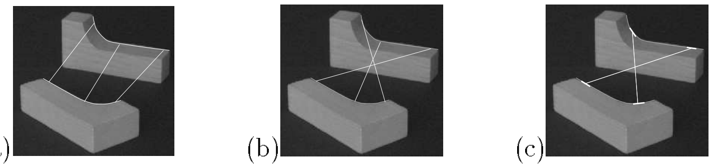

(a) (b) (c)

Fig. 1. The affine transformation bases in (a) and (b) are formed from 3 point correspondences, whereas the basis in (c) is formed from 2 point correspondences and the associated tangents. However only (b) and (c) correctly defines the relation between the two image curves.

$$\boldsymbol{A}_{21} : \boldsymbol{x}_{2j}(t_{2j}) \longmapsto \boldsymbol{x}_{1j}(t_{1j}), \qquad j = 1, .., N \tag{2}$$

where $\boldsymbol{A}_{12}$ is the matrix mapping points $\boldsymbol{x}_{1j}(t_{1j})$ to $\boldsymbol{x}_{2j}(t_{2j})$ and $\boldsymbol{A}_{21}$ is the matrix for the inverse mapping. However, *the matrices which transform $\boldsymbol{x}_{1j}$ to $\boldsymbol{x}_{2j}$ and vice-versa are derived via these points*:

Proposition 1. (Transformation as a function of curve parameters) *A transformation matrix $\boldsymbol{A}_{12}$ which is derived from a transformation basis comprising of N correspondences between two sets of curves and P contour derivatives may be represented by a tensor function M operating on the $2N$-dimensional curve parameter space, ie.*

$$\begin{aligned} M : \mathbb{R}^{2N} &\longrightarrow \mathbb{R}^{2N+P} \\ (\boldsymbol{t}_1, \boldsymbol{t}_2) &\longmapsto \boldsymbol{A}_{12}(\boldsymbol{t}_1, \boldsymbol{t}_2) \end{aligned} \tag{3}$$

where $\boldsymbol{t}_1 = [t_{11}\ t_{12}\ \ldots\ t_{1N}]^T$ is a vector of curve parameters for the first set of curves, and $\boldsymbol{t}_2 = [t_{21}\ t_{22}\ \ldots\ t_{2N}]^T$ the vector of curve parameters for the second set.

The $2N$-dimensional space of curve parameters may be considered a generalisation of Van Gool *et al.* 's [12] 2D *Arc-Length Space* (used for comparing pairs of points on curves) to higher dimensions for establishing multiple correspondence pairs.

Bases for Affine Symmetry. As we previously described in [4], there are three degrees of freedom in an affine symmetry transformation, more intuitively expressed as the location of the symmetry axis (2 DOF), and the angle of skew (1 DOF). We explained how *a single point correspondence together with the tangents at these points form an affine symmetric basis.*

A pair of corresponding points $\boldsymbol{x}_i(t_i) = [x_i(t_i)\ y_i(t_i)]^T,\ i = 1, 2$ is related by:

$$\boldsymbol{x}_2\, \boldsymbol{Z}_2 = \boldsymbol{x}_1\, \boldsymbol{Z}_1 \tag{4}$$

where $\boldsymbol{Z}_1$ and $\boldsymbol{Z}_2$ are matrices containing the symmetry parameters ϕ (the orientation of the axis of symmetry), C (the perpendicular distance of the origin to the axis), and α (the angle of skew). *These parameters are functions of the basis correspondences $\boldsymbol{x}_1(t_1)$ and $\boldsymbol{x}_2(t_2)$ and hence are functions of the arc-length curve parameters t_1 and t_2.* Further details may be found in [3].

2.2 Correspondences Formed with Neighbourhood of Curves

The use of landmark points in establishing transformation bases for curve matching is disadvantageous in that the transformations may be sensitive to errors in localisation of these points. We propose instead an approach based on matching *regions* of curves, in which **(i)** hypothesised corresponding points (which need not be restricted to landmark points) are allowed to vary for limited distances along the contours; **(ii)** the neighbourhood of curves at these points should not only be used to verify the validity of the hypothesised transformation (via the *edge-continuity assumption*), but also be *integrated* into an optimisation process to *localise* the correspondences; and **(iii)** the rest of the curves are gradually incorporated to improve the accuracy should the neighbourhood of the correspondences be well matched after the initial optimisation process, *and* appear *geometrically salient* (defined later in Sect. 2.3).

Localisation of Curve Neighbourhood Correspondences. In order to have a measure of the differences between two sets of curves, we define the *deviation* between curves C_1 and C_2:

Definition 2. (Deviation)
*The **Deviation** D_{12} from curve C_1 to curve C_2 is defined as the integral of squared distances from points on C_1 perpendicular to C_2, ie.*

$$D_{12} = \int_{C_1} \|\boldsymbol{x}_1(\tau) - \boldsymbol{C}_2\|^2 \, d\tau \tag{5}$$

where τ is the arc-length parameterisation for curve C_1, and $x_1(\tau)$ is the position vector for the point with parameter τ on curve C_1. See Fig. 2(a).

Note that this is not invariant to the interchange of curves, ie. $D_{12} \neq D_{21}$, but it is suitable to our analysis as well as for practical purposes.

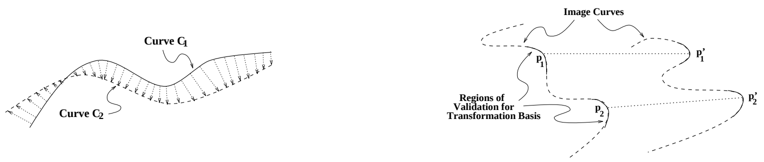

(a) (b)

Fig. 2. (a) The deviation of C_1 to C_2 is the integral of squared distances from points on C_1 perpendicular to C_2, as defined by definition 2. (b) A neighbourhood-validated basis comprising of 2 correspondences is shown.

From (3), it is observed that as basis points slide along the associated curves, the transformation basis and matrix change accordingly. If the basis defines a correct hypothesis for the transformation relating the two sets of curves, the projection of one set of curves via the transformation will match the second set of curves exactly. Since not all parts of the curves are necessarily available in the images, bases will have to be classified according to the *curve neighbourhood support* available at the basis points:

Definition 3. (Neighbourhood-Validated Basis)
A neighbourhood-validated basis, or ***NV-basis****, is a transformation basis which satisfies the requirement that the projection of the curve neighbourhood at all basis points match exactly the curve neighbourhood at their respective corresponding basis points. This is a necessary but insufficient condition for the basis to represent the correct transformation between the two sets of curves. See Fig. 2(b).*

The goal of *localising* the correct set of correspondences to define the transformation between two sets of curves can therefore be formulated as an optimisation problem on the $2N$ dimensional domain of curve parameters. However this set of of correspondences is not unique since any pair of actual corresponding points between the two curves may be used. We instead reformulate the optimisation problem:

Proposition 4. (Correspondence localisation as a minimisation)
Suppose a transformation basis comprising N correspondences is needed to establish the transformation relating two sets of curves. Given N points with the curve parameters $\boldsymbol{t}_1$ on the one set of curves, an NV-basis may be found by recovering the corresponding points with curve parameters $\boldsymbol{t}_2$ on the opposite set of curves through solving the minimisation problem

$$\min_{\boldsymbol{t}_2} \left\{ \sum_{j=1}^{N} \int_{N\{t_{1j}\}} ||\boldsymbol{A}_{12}(\boldsymbol{t}_1, \boldsymbol{t}_2)\, \boldsymbol{x}_{1j}(\tau) - \boldsymbol{C}_{2j}||^2 \, d\tau \right\} \tag{6}$$

where $\boldsymbol{A}_{12}$ is as defined in (3), C_{ij} is the curve on which $\boldsymbol{x}_{ij}$ lies, and $N\{t_{ij}\}$ is the curve neighbourhood of arbitrary size about t_{ij}. The N points on the first set of curves are the fixed basis points (or ***pivot points****), whereas the N corresponding points on the second set of curves are the* ***free*** *basis points.*

Optimisation Process for the Localisation of Correspondences. We outline this *correspondence localisation* process in terms of a *prediction-verification* mechanism as follows:

1. Initially extract from both sets of curves a number of interest points (eg. points with large cornerness [7]).
2. Hypothesise a transformation basis by assuming correspondences between N pairs of points on the two sets of curves and using P contour derivatives at these points,
3. Project the curve neighbourhood at the pivot points via the hypothesised transformation and compute the total deviation between the projected curve neighbourhood and the second set curves;
4. Predict the best locations on the second set of curves for the free basis points by minimising the deviation with an optimisation algorithm operating on the N dimensional space of curve parameters;

5. Cycle through steps 2 and 4 until a minima in deviation is reached, or if the estimation diverges.
6. Determine the *geometric saliency* (Sect. 2.3) and *normalised basis likelihood* (Sect. 2.4) of the hypothesised transformation basis to decide if the basis should be admitted.

Figure 3 demonstrates this optimisation process for an affine symmetry basis.

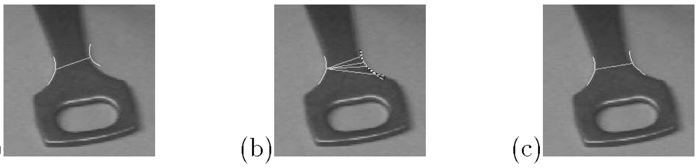

(a) (b) (c)

Fig. 3. (a) An affine symmetric transformation basis representing the initial hypothesis for the transformation relating the pair of image curves; (b) the free basis point is localised by an optimisation process; (c) the final transformation basis is obtained which correctly represents the affine symmetry between the two contours.

2.3 Geometric Saliency of Transformation Bases

In [2], Brady discussed the need in vision problems to work from locations offering the tightest constraints ('seeds') to locations of decreasing constraint. For example in point registration, corners are considered first followed by edges. In the case of geometric curves, curvature is often cited as a measure of its constraint potential [5]. We show that this is *not necessarily so*, since the localisability of correspondences is *both* dependent on shape as well as the transformation class and defining parameters, especially in cases when curvature does not undergo a monotonic mapping.

It may be shown that the localisation of correspondence can be poor even when there is insignificant deviation of projected and actual curve neighbourhood. For an example, consider the case of bilateral symmetry for which a transformation basis comprises of a point-to-point correspondence and no derivatives. In an attempt to establish the transformation relating the two sets of symmetric curves in Fig. 4(a), a wide range of incorrect bases which have strong neighbourhood support may be formed from a pivot point with a free corresponding point – it is impossible to localise the correct basis accurately. A more discriminating transformation basis is shown in Fig. 4(b). Note that the curvatures at the basis points in Fig. 4(b) are smaller than the curvatures in Fig. 4(a), contrary to expectations that points with larger curvatures are more salient.

In order to select the transformation bases with greater discriminatory power, we derive a quantitative definition for the *geometric saliency* of a correct transformation basis.

Derivation of Geometric Saliency. Consider a NV-basis subjected to perturbation of its free basis points. Since the pivot points are fixed, the transformation matrix $\boldsymbol{A}$ relating pivot points to free points in (1) will be rewritten as a function

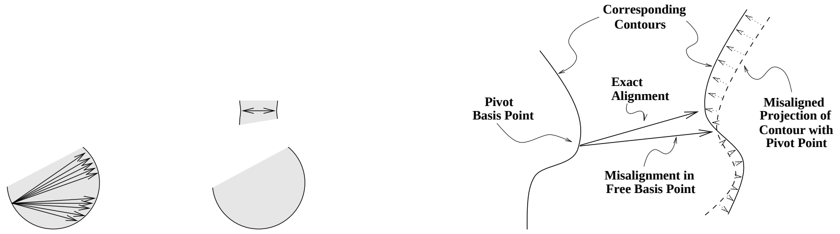

(a) (b) (c)

Fig. 4. In (a), it is impossible to localise the correct bilateral symmetric NV-basis. (b) shows a highly localisable NV-basis which correctly represents the bilateral symmetry. In (c), when a free basis point in an NV-basis is perturbed by a small error, the resultant deviation between the projected and actual curve neighbourhood may be expressed in the form of equation 7.

of only the curve parameters associated with the free points, ie. $\boldsymbol{A}(\boldsymbol{t}_2)$. Additionally, the deviation D_{12} from (5) is defined as a function of the perturbation, and written $D(\boldsymbol{\epsilon})$, where $\boldsymbol{\epsilon} = [\ \epsilon_1\ \epsilon_2\ \ldots\ \epsilon_N\ ]^T$ is a vector of the perturbation in $\boldsymbol{t}_2$. Finally, all curve parameters are assumed to be expressed in terms of arc-length.

For *small* $\boldsymbol{\epsilon}$, the deviation between the projected curve neighbourhood of the pivot points and the curve neighbourhood of the free basis points for the NV-basis may be written as

$$D(\boldsymbol{\epsilon}) = \sum_{j=1}^{N} \int_{a_j}^{b_j} r_j^2 \, d\tau \tag{7}$$

$$r_j = \boldsymbol{x}_{1j}^T(\tau)\ \left(\boldsymbol{A}(\boldsymbol{t}_2) - \boldsymbol{A}(\boldsymbol{t}_2 + \boldsymbol{\epsilon})\right)^T\ \boldsymbol{k}_j(\boldsymbol{t}_2, \tau) \tag{8}$$

where b_j and a_j demarcate the parameter boundaries for the curve neighbourhood of pivot point j, and $\boldsymbol{k}_j$ are the unit normals to the *projected* curve neighbourhood of pivot point j. The curve neighbourhood of the free basis points do not appear in this equation since for a NV-basis they are obtained simply by projecting the curve neighbourhood of the pivot points. See Fig. 4(c).

Taking derivatives,

$$\boldsymbol{\nabla}_{\boldsymbol{\epsilon}} D(\boldsymbol{\epsilon}) = -2 \sum_{j=1}^{N} \int_{a_j}^{b_j} r_j\ \boldsymbol{\nabla}_{\boldsymbol{\epsilon}} \left(\boldsymbol{x_{1j}}^T(\tau) \boldsymbol{A}(\boldsymbol{t_2} + \boldsymbol{\epsilon})^T \boldsymbol{k_j}(\boldsymbol{t_2}, \tau)\right)\ d\tau \tag{9}$$

Since $\boldsymbol{\epsilon} = \boldsymbol{0}$ represents the least-squares solution for (7), $r_j = 0$ and hence $\boldsymbol{\nabla}_{\boldsymbol{\epsilon}} D(\boldsymbol{0}) = 0$. The Hessian matrix $\boldsymbol{H}(\boldsymbol{\epsilon})$ for D at $\boldsymbol{\epsilon} = \boldsymbol{0}$ is non-zero but has a simplified, *positive semi-definite* form:

$$\boldsymbol{H}(\boldsymbol{0}) = 2 \sum_{j=1}^{N} \int_{a_j}^{b_j} \boldsymbol{G}_j\ \boldsymbol{G}_j^T\ d\tau \tag{10}$$

$$\boldsymbol{G}_j = \boldsymbol{\nabla}_{\boldsymbol{\epsilon}} \left(\boldsymbol{x_{1j}}^T(\tau) \boldsymbol{A}(\boldsymbol{t_2})^T \boldsymbol{k_j}(\boldsymbol{t_2}, \tau)\right) \tag{11}$$

where $\boldsymbol{G}$ represents the Jacobian of $r_j(\boldsymbol{\epsilon}, \boldsymbol{t}_2, \tau)$ at $\boldsymbol{\epsilon} = \boldsymbol{0}$. Furthermore since $\boldsymbol{H}(\boldsymbol{0})$

is *positive semi-definite* from (10), all its eigenvalues are real. We therefore define our *geometric saliency* as follows:

Definition 5. (Geometric Saliency)
Given a NV-basis for which $\boldsymbol{H}(\boldsymbol{\epsilon})$ is the Hessian function of the deviation $D(\boldsymbol{\epsilon})$ as defined in (7), the ***Geometric Saliency*** *S of this basis is defined as the* ***smallest eigenvalue*** *of $\boldsymbol{H}(\boldsymbol{0})$, ie.*

$$S = \min(\lambda) \tag{12}$$

$$\textit{subject to} \qquad \boldsymbol{H}(\boldsymbol{0})\,\boldsymbol{\epsilon} = \lambda\,\boldsymbol{\epsilon}, \qquad \lambda \geq 0 \tag{13}$$

The associated normalised eigenvector of S, denoted by $\boldsymbol{v}_S$, represents the normalised $\boldsymbol{\epsilon}$ which causes the minimum increase in the deviation $D(\boldsymbol{\epsilon})$.

We noted that the form of the first derivative matrix '$\boldsymbol{M}$' used in the Plessey corner finder [7] is similar to $\boldsymbol{H}(\boldsymbol{0})$ except that it is specifically two-dimensional and applied to the image domain.

The $D(\epsilon)$ energy functional surfaces in 2D curve parameter space is graphically illustrated for various bases representing *similarity* (scaled Euclidean) transformations in Fig. 5. These bases are formed with two point correspondences and no contour derivatives (four degrees of freedom represented by a 2D translation, a rotation and an isotropic scaling). The eigenvectors associated with the geometric saliencies are in the directions of the most gradual ascent at the minimum point. It is seen that the transformation basis in Fig. 5(a) is preferred to the basis in Fig. 5(b) by virtue of its larger geometric saliency and hence greater localisability. This is in spite of the fact that they both represent the correct similarity transformation relating the two curves.

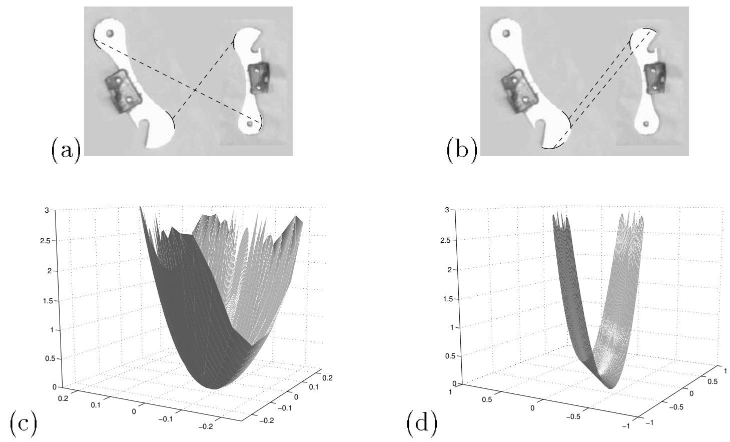

Fig. 5. The similarity (scaled Euclidean) NV-bases shown in (a) and (b) are formed with two point correspondences, and the deviation functionals are shown respectively in (c) and (d). The geometric saliencies are related to the directions of the most gradual ascent.

Geometric Saliency for Affine Symmetry. The Hessian derived in (10) reduces to a single value since only one pair of corresponding points is involved, and is therefore equal to the geometric saliency. Complete details of the derivation may be found in [3]. Figure 6 shows some transformation bases for affine symmetry. The geometric saliency of the transformation basis in Fig. 6(b) is near zero because the ellipse formed by the base of the key is affinely equivalent to a circle, which has infinite symmetry axes.

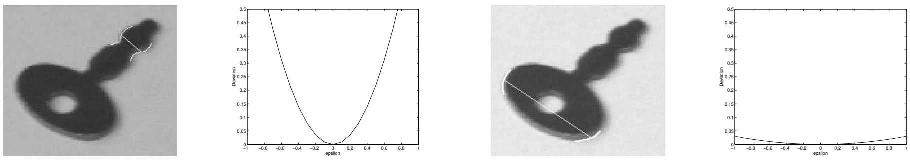

(a) (b) (c) (d)

Fig. 6. The affine symmetry bases in (a) and (c) are formed with one pair of point correspondences. Basis (c) has a near zero geometric saliency since all points on an ellipse are affine-symmetric to each other.

2.4 Statistical Analysis and Normalised Basis Likelihood

When image noise is present, even NV-bases are expected to contain residual deviation in the curve neighbourhood. In this section, we therefore consider the essential task of judging the localisation errors involved when a basis hypothesis is presented.

In a discrete implementation, the minimisation problem given in (6) may be expressed as

$$\min_{\boldsymbol{t}_2} \left\{ \mathcal{D}(\boldsymbol{t}_2) \;=\; \sum_{j=1}^{N} \sum_{i=1}^{Q_j} \boldsymbol{s}_{ji}^T \, \boldsymbol{s}_{ji} \right\} \tag{14}$$

$$\text{where} \qquad \boldsymbol{s}_{ji} \;=\; \boldsymbol{A}(\boldsymbol{t}_2)\, \boldsymbol{x}_{1j}(\tau_{ji}) - \boldsymbol{x}_{2j}(\rho_{ji}) \tag{15}$$

and where Q_j is the number of samples used on the curve neighbourhood of basis point j. ρ_{ji} is determined such that $\boldsymbol{x}_{2j}(\rho_{ji})$ is the nearest point on curve C_{2j} to $\boldsymbol{x}_{1j}(\tau_{ji})$.

It may be shown (please refer to [3]) that by linearising about the solution $\mathcal{D}(\boldsymbol{t}_2) = 0$ we have

$$\mathcal{D}(\boldsymbol{t}_2 + \boldsymbol{\epsilon}) \;=\; D(\boldsymbol{\epsilon}) \;=\; \boldsymbol{\epsilon}^T \, \boldsymbol{H}(\boldsymbol{0}) \, \boldsymbol{\epsilon} \tag{16}$$

In the presence of image noise, errors are introduced into the terms $\boldsymbol{x}_{ij}$, which distort the relation between $\mathcal{D}$ and D. Solving (14) produces the *maximum-likelihood* (ML) estimate $\hat{\boldsymbol{t}}_2$ for the actual $\boldsymbol{t}_2$. The residual deviation $\mathcal{D}(\hat{\boldsymbol{t}}_2)$ arises from a combination of errors in the measurement of s_{ij}'s and the error in the estimation of $\boldsymbol{t}_2$. From the relation $\hat{\boldsymbol{t}}_2 = \boldsymbol{t}_2 + \delta \boldsymbol{t}_2$, $\delta \boldsymbol{t}_2$ represents the *misalignment* error. The error $\boldsymbol{\delta t}_2$ is dominant along $\boldsymbol{v}_S$, the eigenvector related to the geometric saliency.

Since the distribution of noise is not known, it is useful to gauge the accuracy of $\hat{\boldsymbol{t}}_2$ by comparing $\mathcal{D}(\hat{\boldsymbol{t}}_2)$ with the deviation $D(\boldsymbol{\delta t})$ which will arise solely through a misalignment error. Intuitively, we would expect that NV-bases with larger geometric saliencies to be better localised despite having a greater amount of residual deviation. This is easily understood from Figs. 7(a,b), where the deviation functionals shown in Figs. 6(b,d) are subjected to independent errors in the measurement of s_{ij}.

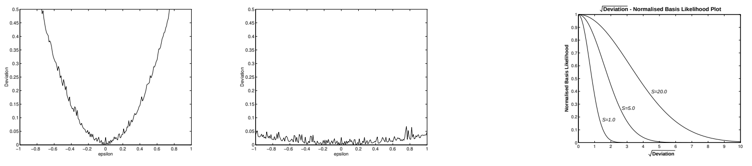

(a) (b) (c)

Fig. 7. The deviation functionals in (a) and (b) are possible scenarios in which the deviation functionals from Fig. 6(b,d) are distorted by errors in the measurement of the image contours. Figure (c) shows the $\sqrt{\text{Deviation}}$ – Normalised-Basis-Likelihood plots for geometric saliencies of 1, 5 and 20.

We make the assumption that $\boldsymbol{\epsilon} = \boldsymbol{\delta t}_2$ is Gaussian-distributed. Furthermore, instead of relating $\boldsymbol{\epsilon}$ to the probability density distribution of $D(\boldsymbol{\epsilon})$ which is highly non-trivial, we relate the deviation $D(q\,\boldsymbol{v}_S)$ to the error q along $\boldsymbol{v}_S$ – the dominant direction of $\boldsymbol{\epsilon}$, ie. assume that $\boldsymbol{\epsilon} = q\,\boldsymbol{v}_S$.

Let

$$z = q\,\sqrt{\boldsymbol{v}_S^T\,\boldsymbol{H}(\boldsymbol{0})\,\boldsymbol{v}_S} \tag{17}$$

Then
$$\begin{aligned} E\{D(q\,\boldsymbol{v}_S)\} &= Var\{z\} \;=\; E\{z^2\} \\ &= Var\{q\}\,\boldsymbol{v}_S^T\boldsymbol{H}(\boldsymbol{0})\boldsymbol{v}_S \;=\; Var\{q\}\,S \end{aligned} \tag{18}$$

Since q is assumed to be Gaussian, so is z. Hence

$$p(z) \;=\; \frac{1}{\sigma_z\sqrt{2\pi}}\exp\left(-\frac{z^2}{2\sigma_z^2}\right) \;=\; \frac{1}{\sigma_q\sqrt{2\pi S}}\exp\left(-\frac{z^2}{2S\sigma_q^2}\right) \tag{19}$$

where $\sigma_z = Var(z)$ and $\sigma_q = Var(q)$. We can therefore define a *normalised basis likelihood* value based on (19):

Definition 6. (Normalised Basis Likelihood)
Given a hypothesis for an NV-basis with geometric saliency S and residual deviation $\mathcal{D}$, the normalised basis likelihood L is defined from (19) by assigning σ_q to be unity and assuming likelihoods to be equal when $\mathcal{D} = 0$:

$$L = \exp\left(-\frac{\mathcal{D}}{2S}\right) \tag{20}$$

L lies between 0 and 1, with $L = 1$ when $\mathcal{D} = 0$. See Fig. 7(c).

From Fig. 7(c) it is apparent that with increasing geometric saliency a greater amount of residual deviation is allowed without a decrease in the likelihood.

2.5 Constraint Propagation and the Refinement of Correspondences

Further refinement of hypotheses may be achieved through a cyclic constraint propagation process:

1. **Curve Support Recovery**. Starting from pivot points, validation checks are propagated along the pivot curves. The goal is to recover the maximum interval of curve correspondences which are in consensus with the transformation hypothesis.

2. **Transformation Refinement**. The restriction that basis points must lie on the image curves is lifted. The optimal transformation relating all corresponding curve intervals is estimated in a least-squares sense through the use of pairs of geometrically-coupled active contours[1]. If the overall deviation between the snakes and the image curves is small, the snakes are used to substitute for the original intervals of image curves, and the algorithm returns to the support recovery step.

Figure 8 shows this process in action.

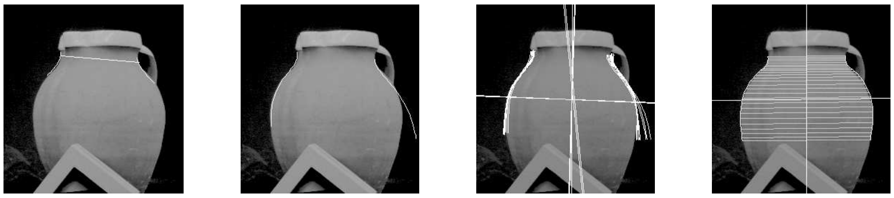

(a) (b) (c) (d)

Fig. 8. In (a), an NV-basis is hypothesised. The curve support recovery proceeds as shown in (b). (c) shows the refinement of the transformation estimate via geometrically-coupled affine-symmetric active contours, with the eventual result obtained in (d). In practice, the process is repeated a number of times.

3 Implementation of a Grouping Process for Affine Symmetric Contours

The grouping of skew-symmetric contours are carried out in the following way:

1. Edgel chains are initially extracted via the Canny edge-detector, to which B-splines are fitted.

2. A number of interest points on curves are extracted and basis hypotheses are made between these points and the *regions* around them, and localised via the process described in Sect. 2.2.

[1] A pair of geometrically-coupled active contours is treated as a single active contour operating in a canonical frame together with a set of transformation parameters mapping the canonical frame to the two image frames. This work will be published in due course.

3. Basis hypotheses are selected if empirical thresholds for geometric saliency and normalised basis likelihood are exceeded.
4. The maximum amount of connected curve support is gathered and the transformation parameters iteratively refined as described in Sect. 2.5.

Note that this is a preliminary implementation. No additional grouping of sets of curves with similar transformation parameters is performed as yet. This is expected to be implemented in due course.

4 Preliminary Results

The affine-symmetric contour-grouping algorithm is applied to the images shown in Fig. 9. B-splines have been fitted to the contour chains obtained from the Canny edge detector, and a number of interest points are selected on these splines.

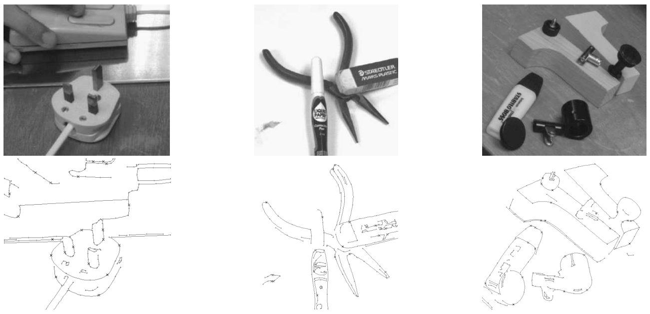

Fig. 9. A number of test images are shown in the top row, and splines fitted to the detected edges are shown in the bottom row. The extracted interest points are marked on the contours with crosses.

Figures 10(a,b,c) show NV-bases which have high normalised basis likelihoods and large geometric saliencies. Further trimming of these 'seeds' is possible by considering the amount of curve support available outside the basis curve neighbourhood. The initial extent of curve support obtained is shown in Figs. 10(d,e,f).

After the application of curve-support / transformation-refinement cycle, we obtain the final results shown in Fig. 11.

5 Conclusions and Future Work

The main aim of this research is to avoid the use of landmark points or high-order differential invariants when establishing correspondences. In the framework of geometrical curves, transformations defined by transformation bases are functions of the curve parameters. Point correspondences are not restricted to matches between isolated point sets, but formed as initial broad matches between intervals of curves. Transformation bases represent hypotheses for the transformation

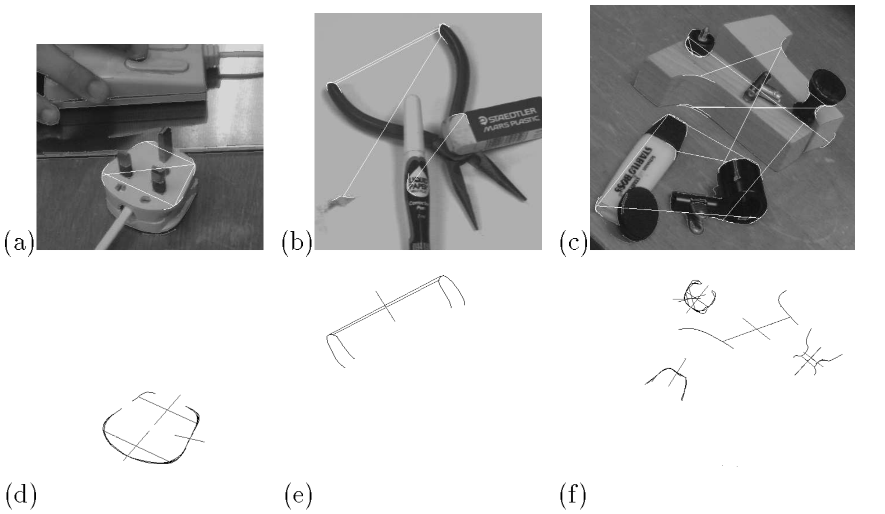

Fig. 10. (a,b,c) show the affine-symmetric NV-bases found from the image contours with have large normalised basis likelihoods and geometric saliencies. Further trimming is achieved by considering the amount of curve support obtained – The strong bases and curve support are shown in (d,e,f) together with the symmetry axes and angles of skew. Different valid hypotheses for the same contours exists in (f).

Fig. 11. The final results obtained after iterating through curve-support recovery and transformation refinement cycles. Corresponding points are marked as shown.

relating curves and may be localised through an optimisation process, terminating in either hypothesis rejection or recovery of neighbourhood-validated bases. We have shown that NV-bases are localisable to different degrees, giving rise to the notion of geometric saliency. In the presence of noise, the accuracy of the estimated correspondences is dependent on the geometric saliency since a highly salient basis is more robust to misalignment errors – the normalised basis likelihood formalises this accuracy. A method is proposed to further validate the hypotheses by recovering additional curve support, and to refine transformation estimates through the use of geometrically-coupled active contours. Preliminary results obtained from applying this theory to the grouping of affine symmetric contours are shown. In the immediate future, further grouping of geometrically-related sets of curves will be implemented.

Acknowledgements. Our appreciation goes to Dr. Andrew Zisserman for providing much valuable feedback. We would also like to thank Jonathan Lawn, Jun Sato and Gabriel Hamid for help and useful discussions. Tat-Jen Cham is supported on an Overseas Research Scholarship and a Cambridge Commonwealth Trust Bursary.

References

1. D.H. Ballard and C.M. Brown. *Computer Vision.* Prentice– Hall, New Jersey, 1982.
2. J.M. Brady. Seeds of perception. In *Proc. 3rd Alvey Vision Conf.*, pages 259–265, Cambridge, Sep 1987.
3. T. J. Cham and R. Cipolla. Geometric saliency of curve correspondences and grouping of symmetric contours. Technical Report CUED/F-INFENG/TR 235, University of Cambridge, Oct 1995.
4. T. J. Cham and R. Cipolla. Symmetry detection through local skewed symmetries. *Image and Vision Computing*, 13(5):439–450, June 1995.
5. M. A. Fischler and H. C. Wolf. Locating perceptually salient points on planar curves. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 16(2):113–129, Feb 1994.
6. R. Glachet, J.T. Lapreste, and M. Dhome. Locating and modelling a flat symmetrical object from a single perspective image. *CVGIP: Image Understanding*, 57(2):219–226, 1993.
7. C. Harris. Geometry from visual motion. In A. Blake and A. Yuille, editors, *Active Vision.* MIT Press, 1992.
8. J. L. Mundy and A. Zisserman. *Geometric Invariance in Computer Vision.* Artificial Intelligence. MIT Press, 1992.
9. P. Saint-Marc, H. Rom, and G. Medioni. B-spline contour representation and symmetry detection. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 15(11):1191–1197, Nov 1993.
10. J. L. Turney, T. N. Mudge, and R. A. Volz. Recognizing partially occluded parts. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 7(4):410–421, July 1985.
11. F. Ulupinar and R. Nevatia. Perception of 3-D surfaces from 2-D contours. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 15(1):3–18, 1993.
12. L. J. Van Gool, T. Moons, E. Pauwels, and A. Oosterlinck. Semi-differential invariants. In J. L. Mundy and A. Zisserman, editors, *Geometric Invariance in Computer Vision*, pages 157–192. MIT Press, 1992.
13. G. A. W. West and P. L. Rosin. Using symmetry, ellipses and perceptual groups for detecting generic surfaces of revolution in 2D images. In *Applications of Artificial Intelligence*, volume 1964, pages 369–379. SPIE, 1993.
14. M. Zerroug and R. Nevatia. Using invariance and quasi-invariance for the segmentation and recovery of curved objects. In J.L. Mundy and A. Zisserman, editors, *Proc. 2nd European-US Workshop on Invariance*, pages 391–410, 1993.
15. A. Zisserman, D. A. Forsyth, J. L. Mundy, and C. A. Rothwell. Recognizing general curved objects efficiently. In J. L. Mundy and A. Zisserman, editors, *Geometric Invariance in Computer Vision*, pages 228–251. MIT Press, 1992.
16. A. Zisserman, J. Mundy, D. Forsyth, J. Liu, N. Pillow, C. Rothwell, and S. Utcke. Class-based grouping in perspective images. In *Proc. 4th Int. Conf. on Computer Vision*, pages 183–188, 1995.

This article was processed using the LaTeX macro package with ECCV'96 style