Skewed Symmetry Detection Through Local Skewed Symmetries

Tat-Jen Cham and Roberto Cipolla
Department of Engineering
University of Cambridge
Cambridge CB2 1PZ, England

Abstract

We explore how global symmetry can be detected prior to segmentation and under noise and occlusion. The definition of local symmetries is extended to affine geometries by considering the tangents and curvatures of local structures, and a quantitative measure of local symmetry known as *symmetricity* is introduced, which is based on Mahalanobis distances from the tangent-curvature states of local structures to the local skewed symmetry state-subspace. These symmetricity values, together with the associated local axes of symmetry, are spatially related in the *local skewed symmetry field (LSSF)*. In the implementation, a fast, local symmetry detection algorithm allows initial hypotheses for the symmetry axis to be generated through the use of a modified Hough transform. This is then improved upon by maximising a global symmetry measure based on accumulated local support in the LSSF — a straight active contour model is used for this purpose. This produces useful estimates for the axis of symmetry *and* the angle of skew in the presence of contour fragmentation, artifacts and occlusion.

1   Introduction

Symmetry, like the other properties described in the Gestalt principles of perceptual grouping [9, 10], represents information redundancy which may be used to overcome noise and occlusion. Under normal viewing conditions we are more likely to encounter skewed symmetry [7], which is planar bilateral symmetry transformed affinely (viewed under *weak perspective* [13]) or perspectively, rather than pure bilateral symmetry.

Friedberg [4, 5] described a method based on moments to detect skewed symmetry axes of 2-D shapes, which assumes prior segmentation of objects. Glachet *et al.* [6] proposed a technique for recovering symmetric axes in line drawings under full perspective projection. Ponce [12] exploited straight-spined Brooks Ribbons and their curvature relationships to pick out symmetrical pairs of points. Mukherjee *et al.* [11] exploited bitangents in objects to define affine bases, from which relative invariants are used to compare pairs of contours with bitangents for symmetry.

Many of these methods apply only to objects which are easy to segment, or are special cases, and do not consider circumstances when symmetry will be useful to perceptual grouping, such as in cases where objects are occluded, or have fragmented contours.

In this paper a local approach to skewed symmetry detection is adopted to handle occlusion. Sensitivity to noise inherent in local measurements is coped with through non-committal classification, uncertainty analysis, and arriving at global decisions by accumulation of independent, local support.

2 Theoretical Framework

2.1 Local Skewed Symmetries

For a point on a curve, its description may be expressed as a Taylor series expansion [3]. The zeroth order term gives the location of the point, the first order term the orientation of the tangent, the second order term the curvature, etc. The locations of two skewed symmetric points provide two constraints on two parameters: the angle of skew and the mid-point between them (the axis of symmetry must pass through this mid-point). Orientation information gives another constraint to fully define the axis of symmetry by the property:

Property 1 *The intersection of tangents to a pair of skewed symmetric contour points lie on the axis of symmetry.*

Having uniquely define the symmetry basis consisting of an axis of symmetry and an angle of skew for the pair of points, we can verify the basis through a further constraint based on curvature:

Proposition 1 *Suppose κ_a and κ_b are the curvatures (inverses of the radii of curvature R_a and R_b shown in figure 1(a)) and σ_a and σ_b are the respective angles between the curve normals and the angle of skew, as shown in figure 1(a). Then if the two points are skewed symmetric, the following must hold:*

$$\frac{\kappa_a}{\kappa_b} = \left(\frac{\cos \sigma_a}{\cos \sigma_b}\right)^3 \tag{1}$$

See [2] for the derivation. The result is similar to that obtained by Ponce [12], but obtained by a different approach. We can extend this constraint verification process to the highest order derivative to which the curve is locally differentiable. This allows us to **extend the definition of local symmetries to different transformation groups**:

Definition 1 *A local symmetry for a particular transformation group refers to a pair of contour points defining a unique symmetry basis within that group. The required number of parameters is obtained through the derivatives of the contours at those points. Additionally, the symmetry basis must also be consistent with at least one other independent constraint provided by the next higher order derivative.*

Notice that the *smoothed local symmetries* proposed by Brady and Asada [1] satisfy this definition, since a symmetry basis under scaled Euclidean transformations may be uniquely defined by the positions of two points; the requirement that the orientations of the tangents to the contours must be equal with respect

to the line joining the points, provides the additional independent constraint. We can also define **local skewed symmetries**, for which the positions and the tangents at the symmetric contour points define a unique symmetry basis under affine transformations, and the curvatures provide the extra constraint to be satisfied in the form of (1).


Figure 1: **(a)**: *Proposition 1 relates the curvatures of a pair symmetric contour points to the angles between the normals and the axis of skew, as shown above; this provides the additional constraint mentioned in definition 1 to define local skewed symmetries.* **(b)**: *In the above state-space the straight line $v = u$ represents the local skewed symmetry subspace. The symmetricity between a pair of contour points is inversely proportional to (1 + the shortest Mahalanobis distance squared from their observed combined state position to the subspace).*

2.2 Symmetricity

Attempting to directly classify pairs of points as local skewed symmetries in the presence of noise and other measurement errors is unlikely to succeed. A better approach would be to adopt a continuous measure of local symmetry. Such a measure can take many different forms, and the one proposed in this paper is based on (1).

In figure 1(b), the state space with state variables $v = \frac{\kappa_a}{\kappa_b}$ and $u = \left(\frac{\cos \sigma_a}{\cos \sigma_b}\right)^3$ is shown for a pair of contour points. The state of the contour-point pair will map onto a single spot $(\mathbf{u_o})$ in this state space, and the straight line $v = u$ is the *subspace for which local skewed symmetry holds*. Noise perturbs the positions of contour-point pairs in this state space, and therefore the proposed symmetry measure is based on the shortest distance to the subspace. Since the error variances are not the same nor necessarily independent for both state variables, then on the assumption that the errors are Gaussian, *Mahalanobis* distance [15] is used instead of the perpendicular (Euclidean) distance.

Assuming that the covariance matrix $\mathbf{K}$ between u and v may be obtained (eg. via automated least-squares B-spline fitting, see [2]), then the Mahalanobis distance squared d^2 is given by

$$d^2 = (\mathbf{u} - \mathbf{u_o})' \, \mathbf{K}^{-1} \, (\mathbf{u} - \mathbf{u_o}) \tag{2}$$

where $\mathbf{u} = [\,u\ v\,]'$ and $\mathbf{u_o}$ is the observed state of the contour-point pair. The shortest Mahalanobis distance to the local skewed symmetry subspace may be

found by substituting $v = u$ into (2) and minimising the quadratic equation. These steps are further elaborated in [2].

This leads us to the following definition for *symmetricity*.

Definition 2 *Given the shortest* Mahalanobis *distance d from the observed state of the pair of contour points a and b, to the local skewed symmetry subspace within the state-space spanned by u and v as described above, the* **symmetricity** *between a and b is given by*

$$\Psi_{a,b} = \frac{1}{1 + d^2} \tag{3}$$

However, work on fitting B-splines automatically is still ongoing, and the current implementation uses *Euclidean* distances instead of Mahalanobis distances.

2.3 The Local Skewed Symmetry Field

In order to *spatially* represent the symmetry evaluation for each pair of contour points in an intuitive and useful way, we define a **local skewed symmetry field** for the set of contours in an image:

Definition 3 *The local skewed symmetry field (LSSF) for a set of contours is formed by:*

- *assigning the symmetricity for a pair of points on the contours to the location of the mid-point of the line joining the pair of points. If more than one pair of points maps to this location, only the pair with highest symmetricity is retained. This is the magnitude of the field at the location.*

- *assigning a direction vector representing the direction of the local axis of symmetry ; This is termed the 'first direction' of the LSSF.*

- *assigning another direction vector parallel to the line joining the pair of points, representing the direction of skew; the 'second direction'.*

(a)　　　　　　　　　(b)　　　　　　　　　(c)

Figure 2: *In order to obtain curvature information to calculate the local skewed symmetry field, cubic B-splines are fitted to the image contours in (a); in the actual implementation, the B-splines are fitted to the output from the Canny edge-detector. The two components of the local skewed symmetry field are represented separately: the symmetricity values are shown in (b), while the first directions (local axes of symmetry) are mapped in (c).*

Figure 2(a) shows a planar symmetric object to which cubic B-splines are manually fitted to the image contours in order to obtain the derivatives for computation

of the local skewed symmetry fields (later implementation differs; details in section 4.2). The magnitude and first direction maps of the local skewed symmetry fields are shown in figure 2(b) and (c).

The ridges in the field bear some similarity to the symmetry set for smoothed local symmetries. The global axis of symmetry appears as a ridge when the data obtained is accurate, and there are also ridges running from the corners. However the ridges of the LSSF are *not* affine-transformed smoothed local symmetries because the SLS are not classified by curvatures.

When errors are introduced, the local, continuous properties of the LSSF allows it to degrade more gracefully under occlusion, fragmentation and noise than if global measures or the strict definition of local skewed symmetries are used. This makes it useful as a tool for recovering the global axis of symmetry from imperfect local skewed symmetries.

3   Symmetry Detection Strategy

In images of objects with complicated contours, it is hard to automatically infer the axis of symmetry purely from its LSSF, especially in the presence of noise and occlusion. A vast amount of computation will be required to analyse and compare different portions of the LSSF. The strategy is therefore to:

1. Obtain initial hypotheses for the axis of symmetry from a fast but approximate symmetry detection process. The symmetry detection process has to use local measures as well to avoid problems of occlusion and incomplete contours.

2. Compute the local skewed symmetry field with the aid of B-splines (to obtain curvature information from the *edge map* of the image).

3. Improve the accuracy of the axis of symmetry and angle of skew by maximising a global symmetry measure from local support in the LSSF.

4   Implementation

4.1   Initial Symmetry Detection Technique

A novel, fast method is introduced for the initial symmetry detection. This does not depend on global measures either and allows local structures to vote independently for the global axis of symmetry and angle of skew, except that the local structures are limited to *low-curvature segments* and *straight edges*.

The symmetry detection process comprises of 5 steps:

1. **Edge-finding and extracting edgel information.**

2. **Extrapolate tangents of edgels in an array.** This is shown in figure 3.

3. **Recover locations of strong intersections.** By recovering only strong intersections, we are effectively considering only the tangents belonging to

Figure 3: *Based on the edge map of the image in (a), the tangents for each edgel are extrapolated in (b). In (c), strong intersection points are recovered which represent intersections of low-curvature segments or straight lines.*

low-curvature segments. The locations of these intersections are more accurate than those from high-curvature segments.

4. **Carry out a reduced Hough Transform on the intersection points.**
The Hough Transform used here parametrise lines by the normal vector from the origin to the line, expressed in polar coordinates (d, θ). Each intersection point maps to a half-sinusoidal curve in Hough space, as may be seen in figure 4. The additional knowledge of property 2 allows us to reduce the Hough Transform by coding only the relevant portions of the sinusoids representing the angular sweep between the directions of the extrapolated tangents:

Property 2 *The orientation of the axis of symmetry must lie between the extrapolation directions of tangents belonging to a skewed symmetric edgel pair.*

This reduces the computation time as well as ignore false hypotheses such as those joining the tips of a serrated edge. The locations of multiple intersections in the Hough Space represent hypotheses for symmetry axis. A number of strong hypotheses are considered.


Figure 4: *The Hough space in (a) shows the result of carrying out a standard Hough Transform on the intersection points shown in figure 3. Knowledge of the directions of tangents at the intersections allows reduction of the complexity of the Hough space (b), from property 2.*

5. **Estimate angle of skew and rank the hypotheses.** This is achieved in a dual-purpose manner by calculating the angle of skew per hypothesis for the axis of symmetry. It is possible to calculate the angle of skew from

only the directions of tangents at *one* intersection point and knowledge of the axis of symmetry, based on (4):

$$\tan\theta = \frac{1}{2}(\cot\gamma - \cot\beta) \tag{4}$$

where θ is the angle of skew with respect to the perpendicular to the axis of symmetry, and β and γ are the angles between the axis of symmetry and the intersecting tangents. Since we will have a number of intersection points contributing votes θ_k for the angle of skew with *different* degrees of accuracy, the optimal estimate is obtained via sensitivity analysis and Lagrange multipliers. The *sample* variances are also obtained, which gives an indication of the accuracy of the estimate and may be used to reject false hypotheses as well. These derivations are given in [2].

4.2 Calculating the Local Skewed Symmetry Field

In this preliminary implementation, hand-fitted cubic B-splines to edge maps are used. A number of sample points per polynomial piece are then extracted with tangent and curvature information. Each sample point is compared with the rest in $\frac{1}{2}N(N-1)$ operations (where N is the number of sample points), to derive the LSSF in three arrays of the same size as the original image, containing separately the symmetricity values, and the two direction data.

4.3 Maximisation of Symmetry Measure in the Local Skewed Symmetry Field

Having obtained a ranked list of hypotheses for the axis of symmetry, we are able to calculate an overall symmetry measure for each of the hypothesis by averaging the **effective symmetricity** values of the points in the LSSF crossed by the axis of symmetry. Taking S as the global symmetry measure and P as the number of sample points, we have

$$S = \sum_{i=1}^{P} \dot{\Psi}_{i\theta} \tag{5}$$

$$\dot{\Psi}_{i\theta} = \Psi_i |\cos(\theta - \alpha_i)| \tag{6}$$

where $\dot{\Psi}_{i\theta}$ is the *effective* symmetricity value at sample point i with respect to the *global* axis direction θ, Ψ_i is the *absolute* symmetricity value measured at sample point i and α_i is the direction of the *local* axis of symmetry in the LSSF at sample point i.

We then proceed to locate the position of maximum symmetry measure for the symmetry axis by initialising a *straight active contour model* [8] (a 'pole'). This 'pole', unlike normal 'snakes', is driven by effective symmetricity values which change according to the orientation of the pole. See figure 5.

The new hypotheses for the axes of symmetry may then be ranked again by their global symmetry measure S. However this does not mean we reject all but one of the hypotheses, since different symmetrical objects, or separate symmetrical components, may be present.

Figure 5: *The above figure diagrammatically shows a straight active contour model (a 'pole') being used to locate the position in the LSSF where the global symmetry measure is maximised. The 'forces' on the pole are based on effective symmetricity (defined in (6)) which depend on* **both** *the absolute symmetricity as well as the angle between the (first) direction of the field and the orientation of the pole. The pole is initialized at the location of the hypotheses obtained from the initial symmetry detection technique forming the first stage.*

5   Preliminary Results

The initial symmetry detection algorithm was tested on a number of synthetic and real images, with straight and curved contours. The results show good performance by the initial technique, in most cases the location of the symmetry axis and the angle of skew are found to a satisfactory degree. The images in figure 6 show some results obtained after the first stage of the algorithm, with back-projection of the relatively intact edge maps.


Figure 6: *Results obtained from the first symmetry detection stage indicate that the first stage performs well with easily segmented objects. The lower row shows the back-projected edge maps to a canonical frame where the symmetry axes are orthogonal.*

Since there may be errors in the estimated axis and ranking, there is a need to use the LSSF to improve the robustness of such estimations. Besides the easily segmented object shown in figure 2, we have also calculated the LSSF for other scenes of objects for which mild occlusion is present, which results in fragmented contours and unwanted artifacts. One of these are shown in figure 7.

Since accurate results are obtained in general from the initial algorithm, *lower ranked* hypotheses have been used to demonstrate the use of the 'pole' in the maximisation of the global symmetry measure in the LSSF. as mentioned earlier. Some of these examples are shown in figure 8.

Figure 7: *The above figure shows the local skewed symmetry field derived from the B-spline fitted edge map of an object, with the magnitude and direction maps shown respectively third and fourth from the left. These will be used in the last stage of the symmetry detection algorithm.*


Figure 8: *The left column shows the initial hypotheses obtained from the first stage of the algorithm. A lower ranked hypothesis is used in (a) in order to demonstrate this stage of the algorithm. The middle column shows the active 'poles' searching for the positions returning maximum global symmetry measures. The final results are shown in the right column. Notice that in (b) where the object is occluded the algorithm still returned good estimates for the axis of symmetry and the angle of skew.*

Overall, the results obtained for the flatter and better segmented objects are particularly good. For objects which have significant thicknesses (deviation from the 2D affine assumption) and those in scenes which include shadows, occlusion or noise, the results still return useful estimates.

6 Future Work

We have presented a new method which does not depend on segmentation of objects or on accurate determination of occlusion-sensitive global measures (eg. moments). Instead, independent support from local structures may be combined to achieve global symmetry detection in the presence of noise and even occlusion, via the local skewed symmetry field. We believe that robust symmetry detection under such circumstances can be exploited to facilitate perceptual organisation as readily as other perceptual groupings such as continuity of contours and T-junctions. Fur-

ther possible applications include integration into a perceptual inference network such as one described in [14], or as an aid in object recognition.

Acknowledgements

References

[1] M. Brady and H. Asada. Smoothed local symmetries and their implementation. *Int. Journal of Robotics Research*, 3(3):36–61, 1984.

[2] T. J. Cham and R. Cipolla. Symmetry detection through local skewed symmetries. Technical Report CUED/F-INFENG/TR 183, Cambridge University, July 1994.

[3] M.P. DoCarmo. *Differential Geometry of Curves and Surfaces*. Prentice-Hall, 1976.

[4] S.A. Friedberg. Symmetry evaluators. Technical Report TR 134, University of Rochester, 1984. (Revised 1986).

[5] S.A. Friedberg. Finding axes of skewed symmetry. *Computer Vision, Graphics and Image Processing*, 34:138–155, 1986.

[6] R. Glachet, J.T. Lapreste, and M. Dhome. Locating and modelling a flat symmetrical object from a single perspective image. *Computer Vision, Graphics and Image Processing – Image Understanding*, 57(2):219–226, 1993.

[7] T. Kanade. Recovery of the three-dimensional shape of an object from a single view. *Artificial Intelligence*, 17:409–460, 1981.

[8] M. Kass, A. Witkin, and D. Terzopoulos. Snakes:active contour models. In *Proc. 1st Int. Conf. on Computer Vision*, pages 259–268, 1987.

[9] D. G. Lowe. *Perceptual Organization and Visual Recognition*. Kluwer Academic Publishers, 1985.

[10] J. D. McCafferty. *Human and Machine Vision: Computing Perceptual Organisation*. Ellis Horwood Series in Digital and Signal Processing. Ellis Horwood, 1990.

[11] D.P. Mukherjee, A. Zisserman, and M. Brady. Shape from symmetry – detecting and exploiting symmetry. To appear in Philosophical Trans. of the Royal Society. (Tech. Rep. OUEL 1988/93).

[12] J Ponce. On characterizing ribbons and finding skewed symmetries. *Computer Vision, Graphics and Image Processing*, 52:328–340, 1990.

[13] L.G. Roberts. Machine perception of three - dimensional solids. In J.T. Tippet, editor, *Optical and Electro-optical Information Processing*. MIT Press, 1965.

[14] S. Sarkar and K. L. Boyer. Integration, inference, and management of spatial information using bayesian networks: Perceptual organization. *IEEE Trans. Pattern Analysis and Machine Intell.*, 15(3):256–273, 1993.

[15] C. W. Therrien. *Decision, Estimation and Classification: An Introduction to Pattern Recognition and Related Topics*. John Wiley, 1989.