

Finding Axes of Skewed Symmetry

Stuart A. Friedberg
Department of Computer Science
University of Rochester, Rochester, NY 14627

TR 127
February 1984

Abstract

A bilaterally symmetric figure on an arbitrarily oriented plane viewed under orthographic projection yields a "skew symmetric" figure whose axes of symmetry and skew constrain the orientation of the underlying plane. Detection of skewed symmetries and extraction of axis orientations is a necessary practical issue that has not received much attention.

We develop a strong necessary condition on the axes of symmetry and skew in a planar skew symmetric figure which constrains them to lie on a one-dimensional locus in a two-dimensional space of orientation pairs.

The practicality of our technique is demonstrated with illustrations from an experimental testbed. A companion paper (TR 134) discusses a number of symmetry evaluators which may be used in conjunction with our strong constraint to determine axes of skewed symmetry.

Key Words: symmetry, skewed symmetry, symmetry constraint.

The preparation of this paper was supported in part by the National Science Foundation under Grants MCS-8302038 and MCS-8209971.

Introduction

Skewed Symmetries

A *symmetric figure* is one that may be formed by reflection of a "half shape" about a line. In Cartesian coordinates, we shall without loss of generality describe the generation of a symmetric figure as discarding all points with negative abscissa and adding the reflection of all points with positive abscissa about the x -axis. The resulting symmetric figure has the same properties (figure or background, intensity, texture, etc.) at points $[x, y]$ and $[x, -y]$.

If we take the points in a symmetric figure and map them (with a shear transformation) to the numerically identical points measured in oblique coordinates, we obtain a *skew symmetric figure*. The figures retain the property of symmetry when coordinates are expressed in terms of the oblique axes and generally do not retain symmetry in terms of the original orthogonal axes. This is a model of orthographic imaging from an arbitrary viewpoint and we sometimes speak of the skew symmetric figure as being in an "image plane" or "image space."

Two axes that together define a set of oblique coordinates in which a figure is symmetric are called *axes of skewed symmetry* for that figure. There may be many such axis pairs for a given figure. The axis corresponding to the axis of reflection is called the *axis of symmetry* and the other is called the *axis of skew*. We shall denote these \mathbf{a} and \mathbf{g} , respectively (Figure 1).

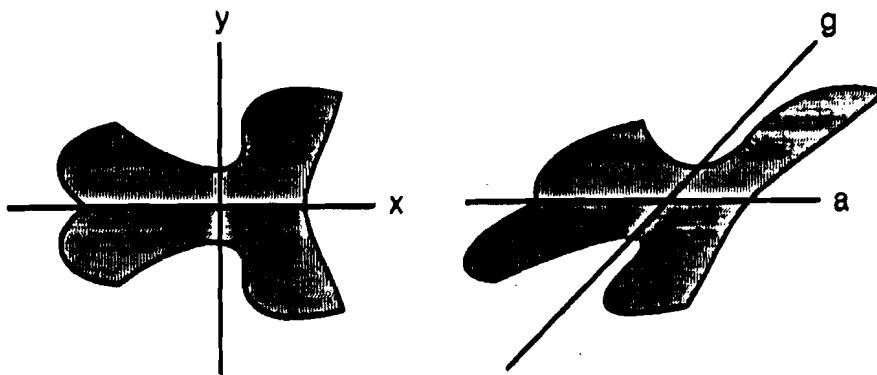


Figure 1. Skew Symmetry under Oblique Coordinate Transformation

Axes of symmetry may not align with the x -axis we use for measurement. *Rotation* of a skew symmetric figure is measured between the axis of symmetry and the x -axis. *Skew* of a skew symmetric figure is measured between the axes of skewed symmetry. Thus, an unrotated, unskewed symmetric figure has a rotation of 0 and a

skew of $\pi/2$. In what follows, the angle of rotation is denoted α ; the angle the axis of skew makes with the x-axis is denoted γ , and the skew is denoted $\beta = \gamma - \alpha$ (Fig. 2).




Figure 2. Rotation and Skew of a Skew Symmetric Figure

Image Domain and Projections

As implied above, we shall use orthogonal Cartesian coordinates for measurements in image space. At times we shall be interested in the coordinates implied by the axes of skewed symmetry. To distinguish measurements in the two systems, we use $\mathbf{p} = [x, y]$ to denote points in the "original" symmetric figure (using the oblique coordinates) and $\mathbf{q} = [u, v]$ to denote points in the "image" skew symmetric figure (using the orthogonal coordinates). To distinguish points in the same coordinates we use \mathbf{p}_1 , \mathbf{p}_2 , and so forth.

To obtain translation invariance in our analysis, we measure all coordinates relative to the centroid or center of mass of the figure. The centroid is invariant under skew and rotation, and has the useful property that all axes of symmetry must pass through it.

Skew symmetric figures are planar figures in the image space that are generated from planar symmetric figures by either of two processes:

- 1) rotation in three dimensions followed by orthographic projection onto the image plane;
- 2) oblique coordinate transformation followed by rotation in the image plane.

A restricted set of three-dimensional rotations will produce skew symmetry from planar symmetry under *perspective* transformation: A rotated axis of symmetry in the original figure must align with the angle of tilt (maximum z -gradient) in the projection.

In both of these processes, the information about the "real" orientation and shape of the planar symmetric figure is lost. As a result skew symmetric figures are ambiguous because there are many figure-transformation pairs that could generate the same skew symmetric figure. The loss is due to two independent sources: *gradient* and *shear ambiguities*.

A gradient ambiguity exists in the three-dimensional process because there is an infinite number of three-dimensional rotations that project the x -axis and y -axis orthographically to the *same* axes of skewed symmetry (Figure 3). This ambiguity is subject to the constraints illustrated in Figure 4 after [Kanade 79], which we have relabelled to avoid confusion between Kanade's notation and our own. The p - q parameter space is the space of x -gradients and y -gradients, which together determine plane orientation relative to reference axes. The solutions for plane orientation given a skewed symmetry lie on the locus of a hyperbola (G_T, G'_T) in p - q space. The asymptotes of the hyperbola are orthogonal to lines drawn in p - q space at angles α and γ .




Figure 3. Gradient Ambiguity




Figure 4. Constraints on Gradient Ambiguity due to Kanade

Shear ambiguity exists in both processes and corresponds to *multiple* axes of skewed symmetry. The number of axes of skewed symmetry for a figure depends upon the shape of the figure. Shear ambiguity is a generalization of the symmetric ambiguity of figures such as the circle and the square, which have an infinite number and four axes of symmetry, respectively. Under skewed symmetry, previously unambiguous figures, such as triangles with all sides of different length, become shear ambiguous (Figure 5). Figures of finite ambiguity under symmetry, such as the ellipse with two axes of symmetry, can become infinitely shear ambiguous under skew symmetry.

Not all figures are more ambiguous under skew symmetry than under orthogonal symmetry. For example, a square retains its four axes of (skewed) symmetry. No figure, however, is less ambiguous under skew symmetry than under orthogonal symmetry. We do not know of a precise characterisation for the degree of shear ambiguity as a function of shape.




Figure 5. Shear Ambiguity of an Arbitrary Triangle

Since gradient ambiguity in the three-dimensional skew symmetry generation process leads to figures that are not uniquely invertible even under a particular projection of axes of skewed symmetry, we choose the two-dimensional process as our model of skewing. Our results are easily related to the three-dimensional rotation and projection process by the well-known constraint of [Kanade 79].

The form of the skew coordinate transformations is $\mathbf{p} \mathbf{T} = \mathbf{q}$, where

$$\mathbf{T} = \begin{array}{cc} \cos\alpha & \sin\alpha \\ \cos\alpha\cot\beta - \sin\alpha & \sin\alpha\cot\beta + \cos\alpha \end{array}$$

Finding axes of skewed symmetry corresponds to calculating skew coordinate transformations that could have taken an unknown symmetric figure into the known image figure. We must be able to deal with the fact that many shapes have multiple axes of skewed symmetry; our goal is to find all of the axes.

Applications and Interest

Finding axes of skewed symmetry is interesting for a number of reasons. First, it seems to present a basic problem in computational geometry regardless of applications. Second, several researchers have applied axes extracted from two-dimensional projections in determining the actual three-dimensional orientation of a figure [Kanade 81], [Kanade, Kender 80], [Kender 78], [Kender 80], [Stevens 79], [Stevens 80]. Axes of skewed symmetry are a useful special case of such axes of interest. Third, actual application of a skewed symmetry finder produces quantifiable evidence for some unintuitive axes of symmetry (i.e., those humans do not perceive as such). Such findings may shed some light on how biological perception of (skewed) symmetries proceeds.

Analytic Constraint

Solution

We shall assume throughout that the figure we are analyzing is in fact derived from (is a skewed version of) a planar symmetric figure.

The development here arose while searching for properties invariant under skew. The starting point is the standard matrix of moments for the figure:

$$\mathbf{M} = \begin{array}{cc} m_{20} & m_{11} \\ m_{11} & m_{02} \end{array}$$

where $m_{ab} = \sum (x_k^a y_k^b)$, summed over all points $\mathbf{p}_k = [x_k, y_k] \in \text{figure}$.

The matrix of moments for a symmetric figure is a diagonal matrix. The moment m_{11} is necessarily equal to zero for a symmetric figure, since for every point $[x, y]$ adding the quantity xy to the moment, there is another point $[x, -y]$ adding $-xy$ to the moment.

The matrix of moments for the original figure, \mathbf{M} , is the sum $\sum \mathbf{p}_k^T \mathbf{p}_k$. Let \mathbf{N} be the matrix of moments measured in the image, $\sum \mathbf{q}_k^T \mathbf{q}_k$. These are necessarily related by $\mathbf{p}_k \mathbf{T} = \mathbf{q}_k$. Then

$$\mathbf{N} = \sum \mathbf{q}_k^T \mathbf{q}_k = \sum (\mathbf{p}_k \mathbf{T})^T (\mathbf{p}_k \mathbf{T}) = \mathbf{T}^T (\sum \mathbf{p}_k^T \mathbf{p}_k) \mathbf{T} = \mathbf{T}^T \mathbf{M} \mathbf{T}.$$

Knowing that \mathbf{M} must be a diagonal matrix and that \mathbf{T} has the form given in the previous section, we solve for α and γ as follows, where the moments m_{11} , m_{20} and m_{02} are from the matrix measured in the image, \mathbf{N} . This gives us:

$$\begin{aligned} \alpha &= \text{atan}((m_{11} \tan \gamma - m_{02}) / (m_{20} \tan \gamma - m_{11})) \\ \gamma &= \text{atan}((m_{02} \cot \alpha - m_{11}) / (m_{11} \cot \alpha - m_{20})) \end{aligned}$$

Fundamental Symmetry Constraint, C.

Taking α and γ over the interval $[0 .. \pi)$ as a periodic parameter space we can plot the locus of (α, γ) pairs meeting constraint C as a connected curve on a torus. Stated another way, we can now describe α as a function of γ and vice versa. C constrains the shear ambiguity in a manner analogous to Kanade's constraint on the gradient ambiguity. In both cases, a two-dimensional parameter space is reduced to a one-dimensional locus of possible solutions. We desire a parameterization of the constraint C giving α and γ as functions of arc length along the locus but have not yet achieved it.

While all solutions under shear ambiguity are represented by points on our locus, not all points on our locus represent actual solutions. This is in contrast to Kanade's locus in which all points represent actual solutions. The difference is fundamentally due to our use of the equality $m_{11} = 0$, which is a necessary but insufficient property of symmetric figures.

Figure 6 shows loci for ellipses with major to minor axis ratios of one (6a), three (6b) and eight (6c). The major axis is aligned with the x-axis for the figures generating the first three loci. The fourth locus (Figure 6d) is generated from the ellipse with axis ratio three under a rotation of 45 degrees. Figure 7 shows loci for the equilateral triangle under zero (7a), 30 (7b) and 60 degree rotation (7c).





Figure 6a.





Figure 6b.


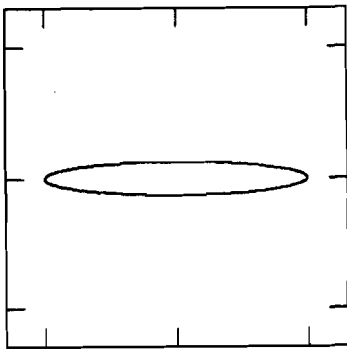


Figure 6c.




Figure 6d.
Constrained Locus of Axes of Skewed Symmetry for the Circle and the Ellipse




Figure 7a.




Figure 7b.




Figure 7c.

Constrained Locus of Axes of Skewed Symmetry for the Equilateral Triangle

Shape Independence

The constraint C is insensitive to absolute shape. The only property of symmetric shapes exploited is the known value of the moment m_{11} . In particular, the constraint is completely insensitive to scaling the original figure along either of the axes of (planar) symmetry. This even applies to nonlinear scaling, such as raising coordinates to a power. This insensitivity to scale implies a strong insensitivity to the shape of the original figure.

Scale the original image in some fashion preserving symmetry. Then the original \mathbf{M} and the new \mathbf{M}' are related by some factors K and L :

$$\mathbf{M} = \begin{matrix} m_{20} & 0 \\ 0 & m_{02} \end{matrix} \quad \mathbf{M}' = \begin{matrix} Km_{20} & 0 \\ 0 & Lm_{02} \end{matrix}$$

This holds even if the scaling was nonlinear. To preserve symmetry, the only restriction is that the scaling along the y -axis (axis of planar skew) must observe $f(-y) = -f(y)$, i.e., be an odd function. The transformation of x coordinates can be completely unrestrained. The actual values of m_{20} and m_{02} from the original matrix \mathbf{M} never appear in the solution.

Need for Evaluation

Since all possible solutions fall on the locus of the graph of α against γ , we have turned a two-dimensional parameter space into a one-dimensional parameter space. Thus the constraint C is a necessary but not sufficient condition to determine the axes of skewed symmetry in a given figure. With a more tractable problem, we

can now apply a variety of search techniques to locate the actual solutions among the possibilities. The strategy is somehow to consider (α, γ) points along the locus until (or as long as) locus points meet some additional condition of symmetry. Several search techniques may be applied: e.g., hierarchical refinement of resolution, local maxima finding and thresholding. A companion paper [Friedberg 84] discusses evaluators of symmetry that can be employed to guide and terminate search.

We do not hold out much hope for additional general analytic constraints due to the nature of the problem. If another constraint provided a different locus in (α, γ) space, the intersection of loci could perhaps determine a finite number of isolated (α, γ) pairs. However, we can construct a figure with any given number of axes of skewed symmetry. Thus any general analytic constraints would have to be of arbitrarily high order to intersect the locus of our constraint an arbitrary number of times. This strongly suggests that we have captured the essence of skewed symmetry in so far as a finite description may be extracted from a figure. That is, apparently *there is no shape number or other description of finite complexity which can fully define skewed symmetry* due to shear ambiguity.

We note further that there are essentially two ways to cope with this if one wants to find actual axes of skewed symmetry. The first approach is to use additional knowledge about the figure which increases in complexity with the complexity of the figure. This might be called the feature-based approach since the useful knowledge is likely to focus on curves, angles, sides, holes, etc. The second approach is to use additional properties about the figure which are bounded in complexity but provide satisfactory results in practice. The feature-based approach has been successfully applied to the detection of lateral and n -fold rotational symmetries in the absence of skew [Bolles 79]. In our experimental testbed we followed the second approach with additional evaluators of "real" symmetry.

Any technique for detecting symmetry in the absence of skew or even in the canonical orthogonal coordinates given by the absence of skew and rotation can be used in conjunction with constraint C to search for axes of skewed symmetry. At the very least, for each candidate (α, γ) the inverse transformation T^{-1} can be applied to all the points in the given figure and the detection technique applied to the result. Feature-based techniques may offer more efficient methods for finding axes of skewed symmetry than search along the locus.

With *a priori* knowledge that the shape of the original figure or the parameters of the imaging process are constrained, additional analytic constraints may be derivable. Such knowledge is often available in specific applications. An example is

determining orientation of parts of known shape in industrial robotics. Prior knowledge of shape may be quite useful in such situations. Even when additional analytic constraints are created, knowledge of original shape may serve to disambiguate instances of shear ambiguity.

Implementation and Experiments

The experimental environment consisted of a Grinnell Systems GMR-26 frame buffer, an Adage RDS-3001 (Ikonas) frame buffer, the Rochester Intelligent Gateway (RIG) and several DEC Vax computers. Software was developed on the Vaxen using the C programming language and locally developed graphics packages to run under the UNIX operating system (Berkeley Standard Distribution, Versions 4.1 through 4.1c). (UNIX is a trademark of AT&T, Vax is a trademark of DEC.)

The experimental testbed used the following procedure. A figure is taken as input data. The mass properties (moments and centroid) of the figure are determined. The pairs (α, γ) lying on the locus of constraint C are calculated to the desired resolution, which is currently 64 equal increments from 0 to π for both α and γ , yielding 128 pairs. A necessary but insufficient evaluator of symmetry is applied to the input figure for each of these (α, γ) pairs. The input figure is displayed; the locus of constraint C is displayed with intensity proportional to the results of the evaluator at that point on the locus, and a plot of the evaluation vs α is displayed. The testbed now permits interactive selection of an arbitrary (α, γ) pair for display of the corresponding axes superimposed on the input figure. The inverse transformation T^{-1} corresponding to the (α, γ) pair can also be applied to the input figure and the results displayed.

Input figures were represented as text for ease of manipulation, where printing characters represent figure and white space represents ground. Procedures to read, write, display graphically, reflect about an axis, and arbitrarily skew and rotate figures were created. The latter two procedures were used to synthesize skew symmetric test figures from known symmetric original figures.

The evaluator used here is the sector symmetry evaluator, E_{SS} , discussed in detail in [Friedberg 84], where the reader is directed for more information. It was chosen from among eight evaluators as most effective for a domain where an area occupancy matrix is used to represent a figure and where significant noise in the figure is anticipated.

As mentioned above, we would like a parameterization of α and γ as a function of arc length along the locus of solutions. This would allow us to take evenly spaced (α, γ) pairs along the locus. The current procedure of taking two sets of pairs, one equally spaced in α and the other equally spaced in γ , is necessary to ensure sufficient resolution along the locus when α is changing rapidly as a function of γ and vice versa.

Experimental Results

In the following figures, taken from the Ikonas display, the upper left shows the input figure. The upper right is a display of the (α, γ) parameter space. The ordinates give α , the abscissa give γ . For both axes the range of values is in $[0 .. \pi)$ periodic. The intensity of each point on the locus shown in the upper right is proportional to the evaluation of the corresponding solution normalized so at least one point is displayed at full intensity. The lower right is a plot of the evaluation (abscissa) against α (ordinate). In this plot the evaluation is normalized by the area of the figure so maximum symmetry is displayed at the top of the scale and minimum symmetry is displayed at the bottom.

The crosshairs superimposed on the upper right quadrant, the lines superimposed on the upper left quadrant and the entire lower left quadrant are related. The crosshairs have selected an (α, γ) pair from the locus of legal solutions. The axes **a** and **g** are then displayed over the figure in the upper left. The lower left shows the input figure "unskewed" to correspond to that candidate solution. The axes superimposed on the lower left figure are the unskewed **a** and **g** which are equivalent to the x- and y-axes. In this photography, the brighter of the two axes (where the distinction can be made) is the axis of symmetry. The dimmer is the axis of skew.

The extreme noise in the figures shown in the lower left quadrant is due to two generations of quantization noise introduced through the skewing process. The first skewing synthesized the figure in the upper left. The candidate solution was then inverted and applied to the upper left figure to produce the lower left figure. There is an additional source of quantization in our experimental testbed: We only resolve solutions to $\pi/64$ radians or about 2.81 degrees. This limits our ability to undo initial transformations precisely.

We earlier noted that the triangle is shear ambiguous in three ways. We begin this section with an example of a scalene triangle. In Figures 8a through 8c the solutions for this example are shown. There are three clear peaks in the plot of E_{ss}

against α , one for each solution. In each case we show the original is a symmetric (isoscalene) triangle. The relative proportions of these solutions vary considerably. A parallelogram is shear ambiguous in four ways. Figures 9a and 9b demonstrate two of the four solutions. The other two solutions are duals of those shown in Figures 9a and 9b, with the axes of symmetry and skew swapped. Notice that two of the solutions are rectangles, while the others are rhombi. Figure 10 shows a figure ambiguous in five ways. Only one of the solutions is shown for brevity. As with the triangle, all the solutions are aligned along a "similar" axis while the relative proportions of length to breadth vary substantially. In Figure 11 we illustrate the infinite shear ambiguity of an ellipse. An arbitrary solution is shown.

Having shown that ellipses provide no information about axes of skewed symmetry, we introduce texture in Figure 12a. The pattern shown was randomly generated with density of about 30%. As displayed it is symmetric about the y-axis. The E_{ss} evaluator is able to find this solution. The pattern is precarious, as shown in Figure 12b. Here the same figure has been skewed and rotated. The solution displayed is the correct one, but it is clearly down in the noise and not distinguished from other possible solutions. We see that a symmetric texture alone will not be enough to pull out axes of skewed symmetry. On the other hand, if the input figure of Figure 12b were sampled from real imagery and not synthesized discretely from discrete originals, *more* of the texture symmetry would be available rather than less and thus the evaluator would work *better*. Figure 12c shows the combined effects for overall shape and a symmetric texture. Here we augment one solution and diminish another. The augmented solution is displayed and is, in fact, the correct solution.

Figure 13a shows the solution for an shape derived from outlines provided by the Production Automation Project at the University of Rochester. There is one solution evaluated by E_{ss} as extremely symmetric. But what are the three lesser but distinct peaks? Figure 13b shows the dual solution for Figure 13a. Figure 13c shows one of the two "diagonal" solutions. These three "extra" local maxima are unintuitive results. An examination of the "original" figures in the lower left shows that there is in fact a "degree" of symmetry present.

While Figure 13 shows plausible but unintuitive axes of skewed symmetry, Figure 14 demonstrates a deficiency of E_{ss} using Stevens's example of the outline of a mousehole. Figures 14a and 14b display the selected "correct" solution, but there are two extremely sharp local maxima of substantial magnitude. Figure 14c selects one of these maxima. While a case could be made for a "minor degree" of symmetry in the original figure displayed in the lower left, an examination of the figure with

the definition of E_{SS} in mind shows that this maxima is likely to be an artifact of the symmetry evaluator [Stevens 80].




Figure 8a.




Figure 8b.




Figure 8c.




Figure 9a.




Figure 9b.




Figure 10.




Figure 11.




Figure 12a.




Figure 12b.




Figure 12c.




Figure 13a.

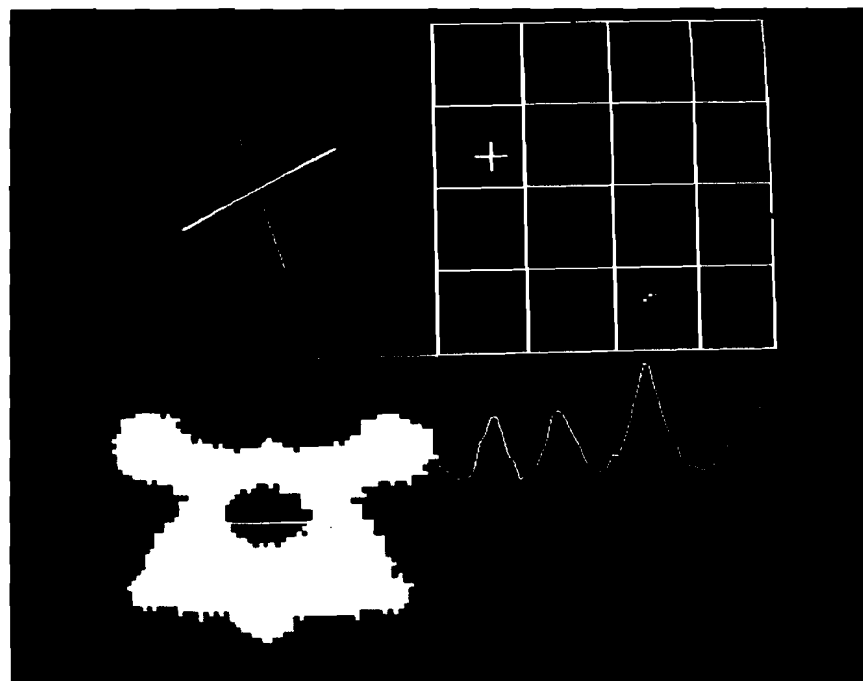


Figure 13b.




Figure 13c.




Figure 14a.




Figure 14b.




Figure 14c.

Unintuitive Results

Frequently the experimental system produced results that were unexpected but quite reasonable upon reflection. Some of the results follow from the nature of the constraint C; others are artifacts of the differences between (our) human evaluations of symmetry and the symmetry evaluators applied to the figures.

A long thin parallelogram presents problems for the detection of axes of skewed symmetry. The behavior at the short ends determines the angle of the skew axis when the symmetry axis is aligned with the long axis of the figure. Many evaluators find it hard to distinguish between various candidate skews. The limiting case is an infinitely long, narrow strip. In this case we can't observe the ends and thus skew is undetermined. This much was intuitive (to us). Experiments later forcibly brought to our attention the symmetry of the locus of the constraint C. An infinitely long, narrow strip has two kinds of ambiguity due to undetermined axes. First, the skew axis is undetermined when the symmetry axis is aligned with the long axis. The only prohibited γ is α . Second, the symmetry axis is undetermined when the skew axis is aligned with the long axis. In this second case the long narrow strip may be thought of as a short, infinitely fat strip and the only prohibited α is γ .

An different kind of unintuitive results developed from experiments with a "Christmas tree" figure (Figure 15). The evaluators initially gave unexpected results from relatively moderate degrees of skew. Analysis of the results showed that the evaluators were finding substantial evidence that the figure was essentially a triangle, well before the amount of skew needed to distort the figure into a triangular blob in (our) human eyes was reached. An objective examination of a skewed Christmas tree indicates that, in fact, the figure differs less and less from a triangle as skew is applied. The difference readily becomes sufficiently small that three solutions for axes of skew symmetry are found while the human eye still sees just one. After sufficient skew distortion we perceive the figure as a mangled triangle rather than a skewed tree, or at least, as some shape which is not a tree.




Figure 15. "Christmas tree" shape

Conclusions

We have addressed the problem of determining axes of skewed symmetry from planar skew symmetric figures. This is of interest in determining three-dimensional orientation of objects. It is a necessary step in order to apply Kanade's constraint on three-dimensional orientation given axes of skewed symmetry that, to our knowledge, has never been successfully developed.

In general there may be many sets of axes for a single figure. We call this *shear ambiguity*. The existence of shear ambiguity indicates that determining three-dimensional shape from a single view *a la* Kanade requires solving a problem with an additional degree of freedom over that discussed in [Kanade 79].

We develop a constraint, C , on the possible sets of axes of skewed symmetry based on the matrix of moments for a figure. Axes must meet this constraint to be axes of skewed symmetry, but meeting the constraint is insufficient to guarantee skew symmetry. General analytic solutions for axes of skewed symmetry are unlikely to exist given the potentially unbounded shear ambiguity. Therefore, evaluation of symmetry for each set of axes on the locus of constraint C is a necessary practical step.

An effective evaluator of symmetry used in conjunction with constraint C has been exhibited.

The practicality of finding axes of skewed symmetry from planar figures has been demonstrated. We give several examples of experiments, ranging from those showing the expected results through those yielding unintuitive results to those demonstrating a limitation of our experimental testbed.

Suggestions for Future Work

A Priori Shape Exploitation

If there are constraints on the original shape of the figure being processed (e.g., figures are known to be portions of automotive connecting rods) or if there are constraints on the axes of skewed symmetry (e.g., aerial photography with known relation between camera and ground coordinates) additional analytic constraints may be developed.

Feature-based Extraction

Figure representations that lend themselves directly to providing constraints on axes of skewed symmetry could be developed that eliminate most if not all of the searching performed by our testbed. Feature-based techniques are quite effective in finding real symmetry. In part this is because they "know where to look" for corresponding features.

Our research is concentrating on polygonal approximations to figures. At a "natural scale" all axes of (skew) symmetry either bisect a side or pass through a vertex of a (skew) symmetric polygon. This places some obvious constraints on candidate values of α . In addition, we are investigating a kind of dual to the (α, γ) parameter space. Constraint C gives a locus in a two parameter space corresponding to the solutions of four equations with five variables: α , γ , three moments in the image figure, and two (unknown) moments in the original figure. Each point (α, γ) along the locus corresponds to a pair of moments (m_{20}, m_{02}) in the original figure. We are interested in looking at the locus in the "moment" space corresponding to constraint C and seeing what happens when a "piece" of the image figure is removed, producing a modified image figure.

Texture Exploitation

Skew transformations significantly alter some forms of texture. Consider radial lines or concentric circles about the centroid as a texture pattern. In an unskewed figure, these lines are of uniform density around the figure. In a skewed figure, the density of these lines varies with position and the axes of skewed symmetry can, in fact, be recovered from the variations in texture density. A generalization of this behavior and its exploitation would provide a useful tool in processing obscured images or in further constraining the axes of skewed symmetry for visible ones. For closely related issues, see [Kanade, Kender 80].

Perspective Generalization

If we adopt an imaging model with perspective projection, we can no longer use constraint C on shear ambiguity or Kanade's constraint on gradient ambiguity. Many tasks in robotics and "near" vision must process images produced under perspective projection. Generalizing *either* constraint would be useful.

Appendix: Calculating Mass Properties (moments)

Throughout our experiments we used an area occupancy matrix representation for our figures. Calculating mass properties with such a representation is trivial. Steve Shafer pointed out to us the utility of perimeter representations in practical computer vision work. This appendix is to indicate how to calculate the moments of an image when an area representation is not available.

The simple technique described here and the others referenced are all variations on Green's theorem and Stokes' theorem, where a double integral is reduced to a single integral. Specifically, an integral over the area of a figure is reduced to an integral around the (directed) perimeter of the figure.

Polygonal Approximation

Assume a polygonal approximation to a figure composed of sequences of connected line segments. Each sequence defines a closed curve where the area to the "left" of the curve is the figure and the area to the "right" of the curve is the ground, i.e., a simple polygon would be described by listing its sides counterclockwise about the center. A hole would be described by listing its sides clockwise about the center.

With a little care a perimeter represented by its component points (as might be produced by a region finder/segmenter) can be treated as a polygonal approximation to a figure, so the simple technique described here can be used without incurring the expense of an explicit line segment approximation.

Polygonal figures with directed perimeters as just described may be broken into a set of triangles. One vertex of each triangle is the (arbitrary) origin and the side opposing that vertex is a line segment of the directed perimeter. The moments calculated for the triangles respect the direction of the line segment opposing the vertex. That is, we want "holes" to remove areas, second moments, etc., while "figures" are to contribute positively. The sums of the moments of these triangles are the moments of the polygonal figure.

We list here simple formulae for calculating various moments on such triangles. Let $[x_1, y_1]$ be the beginning and $[x_2, y_2]$ be the end of the line segment forming part of the directed perimeter of the polygonal figure.

$$m_{00} = (x_1y_2 - x_2y_1)/2$$

$$m_{10} = (y_2x_1^2 + (y_2-y_1)x_1x_2 - y_1x_2^2)/6$$

$$m_{01} = (x_1y_2^2 + (x_1-x_2)y_1y_2 - x_2y_1^2)/6$$

$$m_{11} = (y_2^2(2x_1x_2 + x_1^2) + 2y_1y_2(x_1^2 - x_2^2) - y_1^2(2x_1x_2 + x_2^2))/24$$

$$m_{20} = (y_2x_1^3 + (y_2-y_1)(x_1^2x_2 + x_1x_2^2) - y_1x_2^3)/12$$

$$m_{02} = (x_1y_2^3 + (x_1-x_2)(y_2^2y_1 + y_2y_1^2) - x_2y_1^3)/12$$

$$m_{22} = (y_2^3(x_1^3 + 3x_1^2x_2 + 6x_1x_2^2) + y_2^2y_1(3x_1^3 + x_1^2x_2 - 6x_2^3) - y_1^2y_2(3x_2^3 + 3x_2^2x_1 - 6x_1^3) - y_1^3(x_2^3 + 3x_2^2x_1 + 6x_2x_1^2))/180$$

$$m_{30} = (y_2x_1^4 + (y_2-y_1)(x_1^3x_2 + x_1^2x_2^2 + x_1x_2^3) - y_1x_2^4)/20$$

$$m_{03} = (x_1y_2^4 + (x_1-x_2)(y_2^3y_1 + y_2^2y_1^2 + y_2y_1^3) - x_2y_1^4)/20$$

Chain-Code Representation

[Freeman 74] discusses the calculation of a number of mass properties given the chain code representation of a figure.

Piece-wise Cubic Approximation

[Miles, Tough 83] provide a complete algorithm for computing a piece-wise cubic approximation to a figure and then calculating mass properties for that representation.

Acknowledgements

This work was initially motivated by a programming problem given on a take-home mid-term exam in Chris Brown's Computer Vision course. There were three levels of difficulty: find the rotation angle of a polygonal symmetric figure, find the rotation angle of an arbitrary symmetric figure, and find the rotation and skew angles for an arbitrary skew symmetric figure. A number of enterprising graduate students promptly chose the most difficult (and open) version of the problem. This suggests that making open research problems optional portions of an exam might occasionally be a good idea.

Ronald Loui, David Sher and Stuart Friedberg approached the problem with variations on a boundary pair Hough transform, while Mark Brucks applied ψ -s curves. At that point the Hough transform was sufficient to extract the axes of skewed symmetry without the benefit, or even knowledge, of constraint C. Analytic constraints on parameters α and γ were sought as an afterthought.

Shmuel Tomer assisted us with a partial constraint on the axis of skew. Once the full constraint was discovered, David Sher made many helpful suggestions on refinements and possible further constraints. Their assistance is much appreciated. We would like to thank Mark Brucks for sharing with us his incomplete results on extracting symmetry and skewed symmetry from the ψ -s curves.

Finally, the symbolic computation we went through was tremendously facilitated by the "vaxima" version of the Macsyma program, made available to us by the joint efforts of MIT's Project MAC and University of California-Berkeley's UNIX development and software distribution project.

Bibliography

[Bolles 79]

Bolles, R.C., "Symmetry Analysis of Two-Dimensional Patterns for Computer Vision," *Proceedings, 6th International Joint Conference on Artificial Intelligence*, pp. 70-72, Tokyo, August 1979; also available as Technical Note 186, SRI International, June 1979.

[Freeman 74]

Freeman, H., "Computer Processing of Line-Drawing Images," *Computing Surveys*, 6, 1, pp. 57-97, March 1974.

[Friedberg 84]

Friedberg, S.A., "Symmetry Evaluators," TR-134, Department of Computer Science, University of Rochester, March 1984.

[Kanade 81]

Kanade, T., "Recovery of the Three-Dimensional Shape of an Object from a Single View," *Artificial Intelligence*, 17, pp. 409-460, 1981; also available as CMU-CS-79-153, Computer Science Department, Carnegie-Mellon University, October 1979.

[Kanade, Kender 80]

Kanade, T. and Kender, J.R., "Mapping Image Properties into Shape Constraints: Skewed Symmetry, Affine-Transformable Patterns, and the Shape-from-Texture Paradigm," *Proceedings, IEEE Workshop on Picture Data Description and Management*, pp. 130-135, Asilomar, CA, August 1980; also available as CMU-CS-80-133, Computer Science Department, Carnegie-Mellon University, July 1980.

[Kender 78]

Kender, J.R., "Shape from Texture: A Brief Overview and a New Aggregate Transform," *Proceedings, DARPA Image Understanding Workshop*, pp. 79-84, Carnegie-Mellon University, November 1978.

[Kender 80]

Kender, J.R., "Shape from Texture," Ph.D. Thesis, Computer Science Department, Carnegie-Mellon University, November 1980.

[Miles, Tough 83]

Miles, R.G. and Tough, J.G. "A Method for the Computation of Inertial Properties for General Areas," *Computer-Aided Design*, **15**, 4, pp. 196-200, July 1983.

[Stevens 79]

Stevens, K.A., "Representing and Analyzing Surface Orientation," in P.H. Winston and R.H. Brown (eds). *Artificial Intelligence: An MIT Perspective*, Vol. 2. Cambridge, MA: MIT Press, 1979.

[Stevens 80]

Stevens, K.A., "Surface Perception from Local Analysis of Texture," Technical Report AI-TR-512, AI Laboratory, Massachusetts Institute of Technology, 1980.