Shape from Contour Using Symmetries

Shiu-Yin Kelvin Yuen *
Cognitive Studies Programme,
University of Sussex,
Brighton BN1 9QH
UK
Janet : kelviny@cogs.susx.ac.uk

Abstract

This paper shows that symmetry is essential to shape from contour and indicates problems with existing measures, based on energy and information.

The paper is divided into two parts : The first part establishes the importance of shape from contour using symmetry from a computational theoretic viewpoint. The second part proposes algorithmic solutions to the problem of symmetry finding.

We have omitted all proofs and some parts from this paper due to lack of space. They may be found in [1]. It also contains some more references on shape from contour, reflectional and rotational symmetry. We encourage you to read it.

1 Introduction

1.1 Shape from contour

Shape from contour is the interpretation of a single line drawing as the projection of a three dimensional entity. Following Levitt [22], we model a line drawing or a binary digital image L as a two dimensional point set S. A pixel belongs to S if it is part of L. Then for a line drawing with n pixels,

$$S = \{(c_i, d_i)/i = 1, ..., n\} \tag{1}$$

where (c_i, d_i) is the (known) x-y position of the i^{th}*pixel*. The *shape from contour problem* is to solve for a three dimensional point set S' such that

$$S' = \{(c_i + u_i\, t_x, d_i + u_i\, t_y, u_i\, t_z)/i = 1, ..., n\} \tag{2}$$

*partially supported by the Croucher **Foundation**

where u_i is the unknown parameter and (t_x, t_y, t_z) is the known projection vector (assuming parallel projection), such that some measure of simplicity is extremized.

There have been two streams of research on shape from contour : "top down" and "bottom up".

1.2 Top down vision

Roberts [3], in his seminal work, considered shape from contour as discovering a set of model instances whose projection gives rise to the input line drawing. This research program was extended significantly by Pentland [4, 5]. It consists of three steps. (1) It finds a parameterized model that spans all geometrical objects in the world. (2) It uses a generate-and-test strategy to predict all possible appearances of the objects. (3) It compares the given appearance and finds the best matching model parameters. There are two main difficulties : Firstly, the parameterized models require too many parameters. For example, a superquadric with (restricted) bending and tapering requires 14 parameters ([5], p. 614). Not withstanding this, the superquadric has difficulties representing some simple shapes, eg. the hexagonal prism. More general parameterized models are needed. The hyperquadrics [6], which includes the superquadrics as a special case, needs many more parameters, however. Secondly, complex nonconvex solids should be represented by more than one parameterized model. This gives rise to a "part-whole" problem : Local best model match needs not be global best. (See Mackworth's conterexamples to Roberts' program in [7]). The top down approach does not appear computationally feasible if we have no knowledge about the visual environment.

1.3 Bottom up vision

Marr [8] and others proposed a radically different approach. These researchers envisage that shape from contour may be accomplished without the use of models. Insteads, it is achieved through some primarily data driven process involving only information processing and physical principles. Shape from contour is cast as a quest for finding a *shape metric*. Ideally, the metric should allow us to tell which of any two shapes are more complex. Shape from contour is then the problem of finding the backprojection that extremizes the shape metric. Under this research paradigm, many shape measures have been reported in the past decade. Almost all of them (1) use some form of energy or information as their basis and (2) assume that the given drawing has been "labelled" [1]. One of the main aims of this paper is to show that (1) is misguided. Moreover, it is well known that it is extremely difficult to label a line drawing in general. So far, "labelling" has been successful only in the simple domains (polyhedra and simple curved solids) with stringent assumptions ("perfect" line drawing, etc). This leads one to suspect whether labelling is a neccessary step. As will be shown in the rest of this paper, we may do away with labelling altogether if we use the more naive but intuitive concept of symmetry.

[1]We are not aware of any shape from contour work (including Kanade's) that does not depend on a labelled drawing. Stevens [9] described a method which uses lines of curvature to infer shape. The problem is that the lines are not given. Tsuji and Xu [10] reported a method to construct a net of similiar lines from a drawing. However, their method is heuristic and depends on first labelling the image into "connect" and "occlude" edges. The recent work on symmetries by Ulupinar and Nevatia [11] also implicitly assumed a labelled drawing (for example, see their shared boundary constraint). The works of Witkin [12] and Kanatani [13], however, have the potential of being labelling independent

1.4 Middle out vision

We believe that shape from contour is neither top down nor bottom up. In natural images as well as in line drawings, there are some *invariants*. These invariants have three dimensional implications. Perception of shape from contour proceeds by identifying these invariants. The perception is the result of selecting a backprojection which best *preserves* the invariants. One of the aims of this paper is to show that the (parallel) projection of reflectional and rotational symmetries form invariants that are robust and may be detected in a highly parallel fashion. Moreover, the amount of effort required to detect these invariants (2 and 3 degrees of freedom respectively) are very small compared with that required to detect parameterized models. These invariants are very useful at inferring the three dimensional shape as well as predicting the hidden and occluded parts. In this paper, we shall present the theoretical results in support of our contention.

This theory echos nicely with the direct perception idea of Gibson [14]. The idea of using invariants has been suggested by Ballard, who proposed that "a major function of the perceptual system is to compute collections of invariants at different levels of abstraction" (pg. 89, [15]) and applied it to the recognition of polyhedra.

It is important to realize that this paradigm is not bottom up vision in another guise. The theory does not limit ourselves to general invariants from images. Invariants may also arise from specific models (eg. the relative spatial dispositions of two projected model lines [16]). In fact, recent works on model-based vision has begun to explore these invariants [15-20]. This paper would however concentrate on image invariants, which are more generally applicable. Theorists have long recognized that both bottom up and top down vision are part of a whole theory. Middle out vision attempts to unify both paradigms through the common notion of invariants.

Part I

Symmetries

2 Symmetries and shape from contour

2.1 A symmetry measure

Symmetry is a prolific phenomenon in the world [21]. It may be defined in terms of three linear transformations in n-dimensional Euclidean space E^n : reflection, rotation and translation. Formally, a subset S of E^n is *symmetric* with respect to a linear transformation T if $T(S) = S$. We shall concentrate on reflectional and rotational symmetries. A reflectional symmetry has a reflectional plane, for which the left half space is a mirror image of the right half. A rotational symmetry has an axis of rotation A and an angle of rotation θ. A rotation of θ about A will give an identical figure. Trivially, a rotation of 0^o will produce an identical figure. We shall count this as one rotational symmetry. Let k be the sum of the number of reflectional symmetries and rotational symmetries. Then because a figure has at least one trivial rotational symmetry, $k > 0$. A sphere is reflectionally symmetric about all reflectional planes passing through the center and rotationally symmetric about all axes passing through the center. Hence for a sphere, $k = \infty$.

This measure was originally proposed by Levitt [22].

Human beings perceive a three dimensional object as the putting together of "perceptual parts" [4]. Suppose S' is decomposed into m "parts" $S'_1, ..., S'_m$ ($m \geq 1$). Let $k(S'_i)$ be the sum of the reflectional and rotational invariance in S'_i. We wish to have as few parts as possible. However, the more points S'_i includes, the lower is $k(S'_i)$. A compromise is needed between the number of parts and the k of each part. Levitt [22] achieved this by minimizing

$$E = \frac{m}{\sum_{i=1}^{m} k(S'_i)} \; (m \geq 1) \tag{3}$$

where m is the number of parts. This measure is used by Levitt to decompose an 2-dimensional point set into a few salient subsets. Some good decomposition results have been obtained by Levitt.

In this paper, we shall use (3) to solve shape from contour problems.

2.2 Interpreting skewed symmetry as real symmetry

A "skewed symmetry" is a planar point pattern such that iff (x, y) exists, $(-x, y)$ exists. The x axis is called the "skewed transverse axis" t_k; whilst the y axis is called the "skewed symmetric axis" s_k. If t_k and s_k are not orthogonal, then we have a skewed symmetry. If they are orthogonal, then the skewed symmetry degenerates to a real (reflectional) symmetry, the reflectional axis of which is s_k.

Stevens [9] 's psychological experiments indicated that human beings will interpret a skewed symmetry as a real symmetry. Although it is possible to find counterexamples (for example, we would not interpret a door as a square), this is almost always true. Hence a good measure should likewise interpret skewed symmetry as real symmetry.

Brady and Yuille [24] showed that maximizing their compactness measure will interpret skewed symmetry correctly as real symmetry. Liang and Todhunter [25] proved that minimizing Barrow et. al. 's equality of angle measure [2] will also interpret skewed symmetry correctly.

These works [24, 25] assume that the skewed symmetry is a parallel projection of a *planar* pattern. That is, the underlying transformation is affine. In the following, we shall also assume that the transformation is affine.

We identify a skewed symmetric subset of the point set S of line drawing L as a part (later, we report how such a part may be identified directly from the line drawing). Hence what we need is the general equation (3) applied to a part S'_i, which is

$$E_i = 1/k(S'_i) \tag{4}$$

This means that the three dimensional interpretation of a part S_i maximizes the sum of the number of reflectional and rotational symmetries of S'_i.

Proposition 1 *Minimizing (4) will interpret a skewed symmetric figure as a real symmetric figure if the projection is an affine transformation.*

Outline of the Proof

It can be proved that the real symmetric interpretation will not have fewer rotational (reflectional) symmetries than any other interpretation.

3 Analysis of shape from contour methods

omitted.

4 Shape from contour measures revisited

In this section, the disadvantages of existing shape measures will be discussed. The symmetry measure has none of these disadvantages, but on the contrary has many delightful advantages. Many of them are unique.

The disadvantages of existing measures are :

d1) The shape measures do not suggest a process to extremize them. Though they are quite simple, it is not apparent how to extremize them. Frequently, brute force quantization is resorted to [24]. A good measure, on the contrary, should suggest a process for computing it.

d2) Measures based on angles are very sensitive to small changes in boundary shape.

d3) The extremization of these measures usually gives two global extrema only. (There are two because of Neckar reversal). Other values of the measures do not have physical meaning (see the section 3.2).

d4) Different dimensions require different measures. For example, the three dimensional version of the compactness measure is

$$(volume)^2/(surface\, area)^3 \tag{5}$$

Some other measures have no obvious higher dimensional generalization at all.

d5) The measures are not applicable to point sets composed of isolated points.

The symmetry measure has none of the above disadvantages. On the other hand, it has the following advantages :

The following are also shared by most existing shape measures :

a1) An n-dimensional sphere has $k = \infty$. The measure considers an n-dimensional sphere as the simplest shape.

a2) It is scale invariant.

The following are unique :

a3) It is applicable to point sets as well as line drawings.

a4) It is "dimension invariant". It stays the same for all dimensions. In n dimensions, a reflectional plane is an (n-1)-dimensional hyperplane and an axis of rotation is an (n-2)-dimensional hyperplane. For example, in two dimensions, they are the line of reflection and the center of rotation respectively. In three dimensions, they are the plane of reflection and the axis of rotation.

a5) It is a *process theoretic* measure. It suggests a two level process theory : Level 1 : Look for symmetries in the image. Level 2 : Find the optimal combination of the symmetries found. Symmetry finding will be discussed in Part 2.

a6) Symmetries are global and robust. They are insensitive to noise as well as occlusion.

a7) It does not rate irregular and random figures. This is consistent with the limits of human ability.

a8) Multiple percepts are allowed. Alternative percepts are those with fewer symmetries. They may still turn out to be favourite, particularly if motivated by a priori preference or beliefs. Allowing multiple percepts is very important. For example, it has been shown [26] that human beings may imagine different hidden part completions given the same line drawing. We also believe that a good measure should not only give the ideal best output, but exhibit graceful degradation [27].

Part II

A computational theory of shape from contour using symmetries

5 Introduction

There are two radically different approaches to solving (3) : Feedback and Generate-and-test.

The main idea behind the first approach is to cast the problem as a self organization problem. The spirit of feedback is behind [28, 29].

However, this approach has three serious difficulties : (1) the feedback function may not have physical significance; (2) the feedback function usually introduces some extra coefficients to the system; (3) it is very difficult to find a suitable feedback function that converges to a good equilibrium. It is refreshing to remember that to grow a good crystal, one has to start at near equilibrium conditions.

As we shall show, in the line drawing, there are *invariants* of symmetries that are preserved under projection. They are usually readily apparent (though the computational method to find them is not apparent). An example is skewed symmetry, which is the projection of a reflectional symmetry. One approach to find these invariants is by a generate-and-test strategy. Two important advantages of this strategy are high parallelism and reliability [30]. The apparent difficulty is the combinatorial explosion. For example, naively, to generate all possible reflectional and rotational symmetries requires 3 and 5 degrees of freedom (d.f.) respectively. One of the key results of this paper is showing that they may be done with 2 and 3 d.f. ! In the rest of the paper, we shall explore the potential of the generate-and-test approach.

Our theory is in two levels :

Level 1 : find invariants due to symmetries in the given image;

Level 2 : find the backprojection that provides the best combination of the invariants.

In this paper, efficient algorithms are reported to tackle Level 1.

In the rest of the paper, we shall assume parallel projection.

The next section will deal with two interesting properties of symmetries.

6 Two interesting properties

Proposition 2 *A reflectional symmetry will project as a skewed symmetry or a reflectional symmetry. Moreover, if it is a skewed symmetry, then if and only if the backprojected reflectional symmetry is planar, the skewed symmetric axis will be straight.*

This leads to two interesting corollaries :

Corollary 2.1

If a skewed symmetric figure is to be interpreted as the projection of a reflectional symmetry, then the reflectional symmetry must be planar.

Corollary 2.2

If a reflectional (real) symmetric figure is to be interpreted as the projection of another reflectinal symmetric figure, then the figure need not necessarily be planar. An example is Fig 1. (pointed out to us by Prof. Harry Barrow).

Note : We are using skewed symmetry in a more general sense. Kanade's initial definition of skewed symmetry has a straight axis.

Proposition 3 *A rotationally symmetric figure may have neither skewed symmetric nor reflectional symmetries.*

An example is shown in Fig 2.

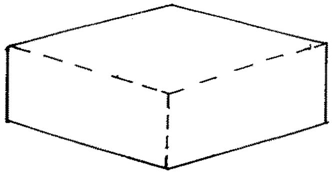

Fig 1

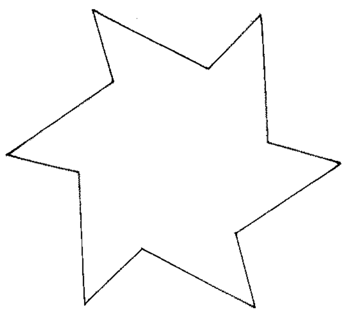

Fig 2

It is well known that a rotational symmetry is a composite of two reflections. However, the last proposition suggests that it is not helpful to look for skewed and reflectional symmetries when finding rotational symmetries.

7 Reflectional symmetry

7.1 Finding skewed symmetry under parallel projection — mapping pair

Levitt [22] proposed an elegant method to detect *reflectional* symmetry in a point set, without requiring a priori that the point set is symmetric. Suppose there are n points. For each of the ${}_nC_2$ pairs, the midpoint is found and assigned a direction perpendicular to the line joining the pair. To this is added the n points with no direction assigned. A Hough transform is then used to find straight reflectional symmetry axes.

We may extend Levitt's method to detect skewed symmetry in the following way. Suppose P_1, P_2 is a pair of points. Then insteads of storing the angle of the perpendicular, we may store the angle β of line $\overline{P_1P_2}$ with some reference axis on a third dimension of the Hough space. High counts on the Hough space correspond to a skewed symmetric axis s_k where the skewed transverse axes are at the same angle with respect to s_k.

This method requires a 3-d Hough space. Below, we describe a variant which requires only a 2-d Hough space.

We point out in proposition 2 that skewed symmetry is the projection of reflectional symmetry. Suppose p', q' are two three dimensional points and p' reflects onto q' (and vice versa). Then we shall say that they are *mapping points* or a *mapping pair.*

Suppose r', s' are another mapping pair. Then clearly $\overline{p'q'}//\overline{r's'}$. Now since a parallel projection is an affine transformation, parallelism is preserved. Hence upon projection $\overline{pq}//\overline{rs}$, where p is the projected point of p' and etc. We note that this is an invariant.

7.2 Extending Kanade's skewed symmetry constraint to parallel projection

We now show how to make use of this invariant to find skewed symmetry in the image and the constraint of the skewed symmetry on the backprojected reflectional symmetry.

Suppose p and q are mapping points. Let the known projection vector be (t_x, t_y, t_z) and σ and τ be the slant and tilt (as defined in [24]) of the reflectional plane . Let

$$\begin{bmatrix} a_x \\ a_y \end{bmatrix} //(q-p) \tag{6}$$

and

$$\begin{bmatrix} a'_x \\ a'_y \end{bmatrix} = \begin{bmatrix} cos\tau & -sin\tau \\ sin\tau & cos\tau \end{bmatrix} \begin{bmatrix} a_x \\ a_y \end{bmatrix} \tag{7}$$

then it can be shown [1] that

$$\sigma = \frac{\pi}{2} - tan^{-1}(\frac{t_z a'_y}{t_x a'_y - t_y a'_x}) \tag{8}$$

(8) may be understood this way. Mapping pairs must be parallel to (a_x, a_y). Hence (a_x, a_y) defines a "mapping direction". Given the mapping direction, we may use (8) to calculate all feasible slant σ and tilt τ of the reflectional plane, assuming a known parallel projection. Once the equation of the reflectional plane is fixed, it is an easy matter to calculate the three dimensional positions of any mapping points.

Recall that proposition 2 states that if the skewed symmetric axis is straight, the reflectional (real) symmetric interpretation must be planar. Hence given a skewed symmetric figure (with straight skewed symmetric axis) in the image plane and a mapping direction, we may apply (8) to find all reflectional planes that will give a symmetric interpretation. We note that 1 d.f. is required for the mapping direction. This is consistent with Friedberg's result [31] that 1 d.f. is needed to detect skewed symmetry. Moreover, the gradient of the reflectional plane has 1 d.f., which is again consistent with the 1 d.f. of Kanade's skewed symmetric heuristic.

However, our formulation is more general because it is applicable to parallel projection, whilst Kanade's formulation assumes orthographic projection. It may also be noted that given τ, the gradient of the reflectional plane is constrained to lie on a line in the gradient space. This may be constrasted with Kanade's heuristic, which specifies that the gradient of the plane *on which the backprojected figure lies on* is on a hyperbola in the gradient space.

Under orthographic projection, the solution is particularly simple, whence it can be shown [1] that the mapping direction is τ.

7.3 Extremum symmetry

We shall define the more general notion of "extremum symmetry" and show that Ulupinar and Nevatias' "parallel symmetry" and "mirror symmetry" (see below) are special cases.

We say that a curve is "mapped" to another curve in the image if a one-one mapping may be defined which maps one curve to another. Of course, the "mapping directions" have to be the same. It is very important to recognize the following :

Lemma 4 *Let c_1, c_2 be image curves and c'_1, c'_2 their backprojection. If a mapping exists which maps c_1 to c_2, a c'_1 and c'_2 may always be found such that c'_1 is reflectionally symmetric to c'_2.*

This lemma indicates that reflecional symmetry may be imposed on almost any two curves. However, this is not so with human beings. Though a mapping may be defined for the two curves in Fig 3, no one would seriously consider that their backprojection is reflectionally symmetric. What extra constraints do we exploit then ?

Ulupinar and Nevatia [11] proposed two special forms of symmetry : "parallel symmetry" and "mirror symmetry". Two analytic curves c_1, c_2 have parallel symmetry if

$$\theta_1(s) = \theta_2(as + b) \tag{9}$$

where $\theta_i(s)$ is the angle the curve made with the positive x axis at curve length s, and a, b are constant. Since a is merely a scaling factor, in the rest of the paper, we shall assume $a = 1$. Then one curve is a translation of the other.

The mirror symmetry is identical to Kanade's skewed symmetry. Recall that Kanade's skewed symmetry has a straight skewed symmetric axis.

Fig 4 shows a reflectional symmetry. The two curves, which are reflective to each other, may be described by

$$(\pm x(t), y(t), z(t)) \tag{10}$$

where

$$c(t) = (0, y(t), z(t)) \tag{11}$$

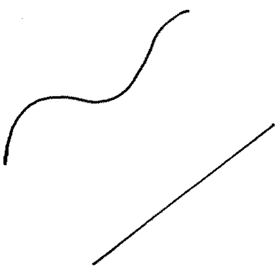

Fig 3

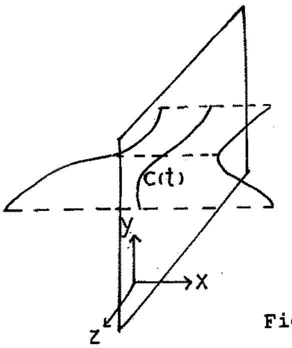

Fig 4

It is not hard to show that

(i) if x(t) = constant, the projected image has parallel symmetry;

(ii) if c(t) is a straight line, the projected image has mirror symmetry

Hence it may be concluded that both parallel and mirror symmetries are special cases of the (parallel) projection of reflectional symmetries.

However, it is incorrect to suppose that parallel and mirror symmetries are the only symmetries that human beings perceive. An example which neither fits parallel nor mirror symmetry is the top face of Fig 5.

We define below the concept of "extremum symmetry", which we believe is the most general form of projected reflectional symmetry that human beings perceive.

Two image curved fragments c_1 and c_2 have "extremum symmetry" if a one-to-one map may be found between points in c_1 and points in c_2 such that

(i) if a point in c_1 is an extremum (curvature maximum, curvature minimum, corner with m incident lines), then the mapping point in c_2 must be an extremum;

(ii) if a point in c_1 is not an extremum, then the mapping point in c_2 is not an extremum.

It is clear that our extremum symmetry subsumes both the parallel and the mirror symmetry. Under this definition, Fig 3 is not a projection of reflectional symmetry whilst Fig 5 is. Once an extremum symmetry is found, the three dimensional position of the backprojected curves may be found using the result of the last subsection.

The concept most closely related to extremum symmetry is Brooks ribbon [32]. A Brooks ribbon is generated by line segments (of varying length) centered on a curved axis. The angle of the line segment is variable. The angle between the line segment and the tangent to the axis is constant. For an extremum symmetry, the angle of the line segment is constant whereas the angle between the line segment and the tangent is variable. If the axis is straight, then both give skewed symmetry with straight axis. But extremum symmetry is in general not Brooks ribbon. Thus we may also say that Brooks ribbon in general does not backproject to reflectional symmetry.

In general, the skewed symmetric axis may be defined as the locus of the midpoints of the mapping pairs of an extremum symmetry. Mirror symmetry restricts the axis to be straight, whilst for extremum symmetry the axis may be an arbitrary connected curve. It is interesting to ask if there are any notions of symmetry intermediate between these two extremes. For example, is "quadric symmetry" — that is, the axis is a quadric curve — possible ? Fig 6 shows an example for which the axis is a circular arc. However, it is not apparant that the two curves have any symmetry relations to a human observer. Indeed, the lower curve has a curvature minimum which is absent in the upper curve, thereby violating the concept of extremum symmetry. We conclude that quadric symmetry is not perceived by human beings.

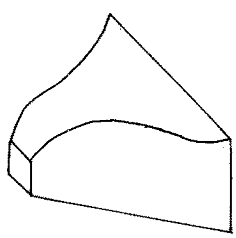

Fig 5

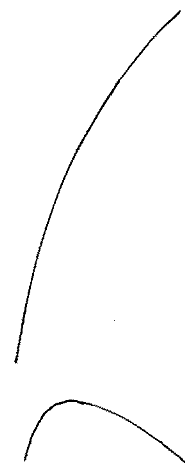

Fig 6

7.4 An algorithm for finding skewed symmetry

It is simple to detect parallel symmetry. Given a mapping, the distance between the two mapping points will remain constant. (One curve is a translation of the other). Hence in this subsection, we shall concentrate on finding an efficient algorithm for detecting skewed symmetry.

Let the mapping direction make an angle α with the positive x axis.

Algorithm 1 (Skewed symmetry finding)

Input : a two dimensional point set.

1. $\alpha \leftarrow 0$;
   ;;; the initial mapping direction is parallel to the x axis

2. clear the midpoint_array;
   rotate the point set by $-\alpha$ about the origin;
   ;;; after the rotation, the mapping direction is parallel to the x axis

3. for all lines L parallel to the x axis do
     if two points p, q lie on the L then
       record the midpoint of p and q on midpoint_array;
     endif;
   endfor;

4. use a Hough transform to find straight lines in midpoint_array
   and store symmetry axes found and their associated point subset;

5. increment α;

6. if $\alpha > \pi$, then exit else goto 2.

Output : a set of symmetry axes and their point subsets.

We have 1 d.f. in α. Finding lines using the Hough transform involves 1 d.f. [1]. Hence the algorithm has two d.f.

The extremum symmetry finding is a special case of the algorithm. The straight line finding is unnecessary. Neighboring midpoints of the mapping pairs of an extremum symmetry may simply be connected together.

It should also be noted that we have not exploited extremum at all in the above algorithm. In actual facts, only α which maps an extremum to another extremum needs be considered. Hence one d.f. may be shrinked. We have not pursued this possibility.

8 Rotational symmetry

8.1 Finding rotational symmetry

Reflectional symmetry does not capture all symmetries in the world. Fig 2, for example, has no global reflectional symmetry. It looks pleasing because it is invariant under rotations of multiples of $\pi/3$. Proposition 3 suggests that rotational symmetry should not be detected by detecting skewed nor reflectional symmetry. This motivates us to study an independent method for detecting rotational symmetry. It should be simple, robust, computationally inexpensive, highly parallel, and should not presuppose a symmetric point set. Last but not least, it should be able to detect rotational symmetry which is not symmetric in the image plane, but is so after a backprojection.

Let us consider the problem of detecting rotational symmetry on the image plane first. Suppose our input is a point set S. Let p, q be two points in S. If p and q are rotationally symmetric, then we may define a rotation $R(o,\theta)$ which rotates p into q. o is center of rotation and θ is the angle of rotation. It is clear that o must lie on the perpendicular of $\overline{pq}$ passing through the midpoint of $\overline{pq}$. Denote this perpendicular as $\overline{pq}'$. Now suppose we have many pairs of points which are rotationally symmetric with the *same* $R(o,\theta)$. Then the perpendiculars must intersect at the origin $o(x_0, y_0)$.

Use the x-y image plane as an accumulator array. If we accumulate all points on $\overline{pq}'$ for every pair p, q on the x-y accumulator array, then there will be a high count at the common origin $o(x_0, y_0)$ (Fig 7).

This transform converts detection of *concentricity* to detection of high counts in the x-y accumulator array. But a rotational symmetry is concentricity and *equality of subtended angles.* That is, the angle subtended by p and q on o, $\angle poq$, must be equal to other angles subtended. Clearly, this common angle is the angle of rotation θ.

One solution is to accumulate in a three dimensional accumulator array (x,y,θ). However, there is a better way to do it.

Scott et. al. [33] reported a brilliant algorithm for detecting "smoothed local symmetries" (SLS). The problem is to find the locus of the SLS axis, points of which are the centers of bitangent circles. Scott used the following analogy. Each point is considered as a small stone dropped on a pond which sets off a one-off ripple. The SLS axis are those points where the wavefronts meet !

We may use the wave propagation analogy to detect rotational symmetry, except that a "wavefront" would be replaced by a "particle" (as dictated by quantum mechanics !). Look at Fig 8. The "particle" is initially at m_0, the midpoint of p and q. The "particle" then splits into two and moves away from m_0 in two directions along $\overline{pq}'$. At "time" $t = \theta$, the "particle" is at m_θ. The position of m_θ may be readily calculated using the right-angled triangle $\Delta m_\theta m_0 p$. If there is a rotational symmetry with angle of rotation θ at o, then many "particles" will meet there and it will have a high count at that instant.

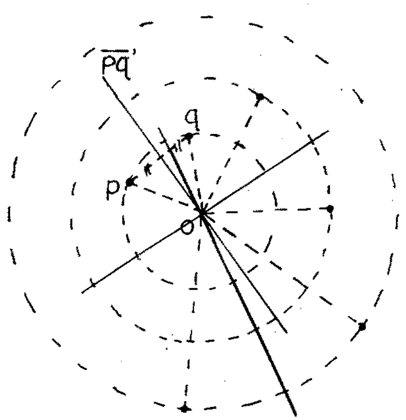

Fig 7

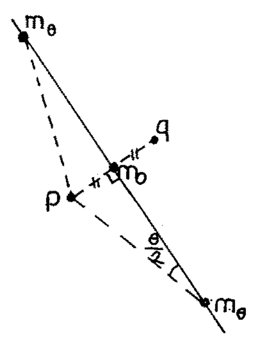

Fig 8

It is clear that this algorithm involves only 1 d.f. (θ) and the space complexity is the x-y plane, which is two dimensional insteads of three dimensional as with the Hough transform.

However, there may exist some rotational symmetries on a slanted and tilted plane. If so, the projection need not have any rotational symmetry. To detect such symmetries, we first backproject the image point set onto a slanted and tilted plane, then perform the symmetry detection above *on that plane.* To specify the slant and tilt entails 2 d.f. Hence our algorithm has a total of 3 d.f. (slant, tilt, θ). Note that the naive way to specify a rotational symmetry axis and its angle of rotation needs 5 d.f. (2 for the x-y projection of the axis, 2 for the x-z projection of the axis, 1 for the angle of rotation). Hence our method represents a significant improvement. It should also be noted that 3 d.f. is well within the computing cabability of today's serial computer. This is even more so with a parallel computer. It is possible that "extremum" may shrink the d.f. further, as with reflectional symmetries. We have not explored this possibility.

So far, our rotational symmetries must lie on planes, which may be slanted and tilted. How about subsets of points which have rotational symmetries but are not necessarily planar ? We note that the normal vector of the slanted and tilted plane is parallel to the rotational axis. Now suppose for such a plane P', there exists two centers of rotation c_1', c_2' with the same angle of rotation, whose projections are c_1, c_2 respectively. Let S_1' and S_2' be the associated three dimensional point sets (which lie on P'). Let S_2'' be a translation of S_2' along the projection vector. Now if $\overline{c_1c_2}$ on the image plane I is parallel to the gradient vector g of plane P' (Fig 9), a rotational axis exists which S_1' and S_2'' are both rotational symmetrical with.

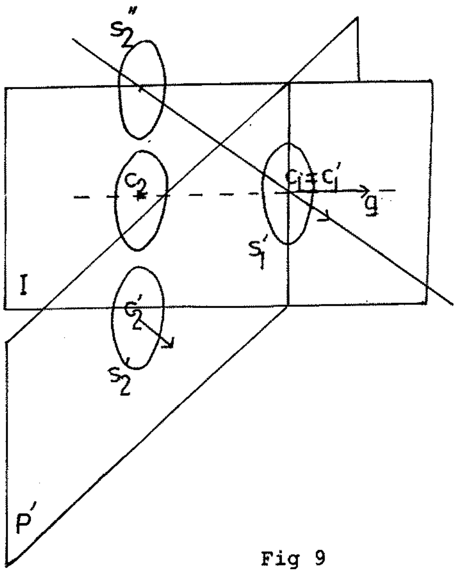

Fig 9

8.2 An algorithm for finding projected rotational symmetry

Algorithm 2 (Projected rotational symmetry finding)

Input : a two dimensional point set; the projection vector

for each slant and tilt do

1. backproject input point set to slanted and tilted plane passing through origin;

2. for each angle of rotation θ do
   for each pair of backprojected point do
      compute m_θ;
      project m_θ to image plane and accumulate
      on the image plane;
   endfor;
   record high points on the image plane, θ, slant and tilt;
endfor;

endfor;

Output : a set of projected rotational symmetries.

In the algorithm, m_θ is projected back to the image plane. This allows us to use the image plane as the accumulator array over and over without worrying about re-quantization.

9 A computational solution for the combination of symmetries found

omitted.

Conclusions

We propose a middle out paradigm for shape from contour problem, which is composed of two stages : Level 1 : find invariants in the given images. Level 2 : find the backprojection that provides the best combination of the invariants. We argue that bottom up vision, which uses low level primitives, does not capture the essential properties of the world. On the other hand, top down vision, which uses full models, are computationally impossible. We believe that some intermediate approach like our middle out paradigm captures the best of both paradigms, yet has none of their attendant disadvantages. It may be regarded as a describe-and-match approach.

We have provided theoretical as well as experimental evidence that suggests that energy/information content based approach to shape from contour is a problematic research program. On the other hand, we are able to show that symmetry has a lot of nice and unique advantages, including the interpretation of skewed symmetry as reflectional symmetry.

Next, we show that symmetry may be implemented with highly parallel, robust and computationally inexpensive algorithms (compared with models and the iterative minimization needed for existing shape measures). We report two such algorithms which also substantially extend the state of the art in skewed symmetry finding and rotational symmetry detection. We have also extended Kanade's skewed symmetry heuristic to parallel projection. We propose a general concept for reflectional symmetry known as "extremum symmetry", which we believe is the most general form of invariance due to reflectional symmetry that human beings perceive. All of our results in this paper are applicable to parallel projection.

This paper is also a step towards a theory of shape [34].

Acknowledgement

First and foremost I wish to thank for the supervision of Dr. David Hogg and for his very careful reading of a draft of this paper. I am indebted to the following people for helpful discussions : Prof. Harry Barrow, Dr. David Hogg, Dr. Hlavac Vaclav, Dr David Young. My research at the Dept. of Electronic Engg, Hong Kong Polytechnic laid the seeds of the present work. Prof. Aaron Sloman's lectures on "Philosophical Foundations on Computing and AI" was helpful in shaping some of my thoughts. I thank Alistair, David and Hilary for help with the SUN workstation.

References

1. Yuen, S.Y., Shape from contour using symmetries, CSRP 141, School of Cognitive and Computing Sciences, University of Sussex (1989).

2. Barrow, H.G. and Tenenbaum, J.M., Interpreting line drawings as three dimensional surfaces, Artificial Intelligence **17** (1981) 75-116.

3. Roberts, L.G., Machine perception of three-dimensional solids, Optical and Electro-Optical Information Processing (1965) 159-197.

4. Pentland, A.P., Perceptual organization and the representation of natural form, Artificial Intelligence **28** (1986) 293-331.

5. Pentland, A.P., Recognition by parts, Proceedings 1^{st} International Conference on Computer Vision, London (1987) 612-620.

6. Hanson, A.J., Hyperquadrics : smoothly deformable shapes with convex polyhedral bounds, Comput. Vision Graph. Image Process. **44** (1988) 191-210.

7. Mackworth, A.K., How to see a simple world : an exegesis of some computer programs for scene analysis, Machine Intelligence **8** (1977) 510-537.

8. Marr, D., Vision (Freeman and Company, New York, 1982).

9. Stevens, K.A., The visual interpretations of surface contours, Artificial Intelligence **17** (1981) 47-73.

10. Tsuji, S. and Xu, G., Inferring surfaces from boundaries, Proceedings 1^{st} International Conference on Computer Vision (1987) 716-720.

11. Ulupinar, F. and Nevatia, R., Using symmetries for analysis of shape from contour, Proceedings 2^{nd} International Conference on Computer Vision (1988) 414-426.

12. Witkin, A.P., Recovering suface shape and orientation from texture, Artificial Intelligence **17** (1981) 17-45.

13. Kanatani, K., Detection of surface orientation and motion from texture by a stereological technique, Artificial Intelligence **23** (1984) 213-237.

14. Gibson, E.J., Constrasting emphasis in gestalt theory, information processing, and the ecological approach to perception, in: J. Beck (Ed.), Organization and Representation in Perception (Lawrence Erlbaum Associates, 1982) 159-166.

15. Ballard, D.H., Form perception using transformation networks : polyhedra, in: C. Brown (Ed.), Advances in Computer Vision vol. **2** (Lawrence Erlbaum Associates, 1988) 79-120.

16. Bray, A., Tracking objects using image disparities, Proceedings 5^{th} Alvey Vision Conference (1989) 79-84.

17. Goad, C., Special purpose automatic programming for 3d model-based vision, Proceedings Image Understanding Workshop (1983) 94-104.

18. Weiss, I., Projective invariants of shapes, Proceedings Image Understanding Workshop (1988) 1125-1134.

19. Rao, K. and Nevatia, R., Generalized cone descriptions from sparse 3-D data, Proceedings IEEE Conference on Computer Vision and Pattern Recognition (1986) 256-263.

20. Ponce, J., Chelberg, D. and Mann, W., Invariant properties of the projections of straight homogeneous generalized cylinders, Proceedings 1^{st} International Conference on Computer Vision, London (1987) 631-635.

21. Rosen, J., Symmetry Discovered (Cambridge Univ. Press, 1975).

22. Levitt, T.S., Domain independent object description and decomposition, Proceedings AAAI-84 (1984) 207-211.

23. Kanade, T., Recovery of the three-dimensional shape of an object from a single view, Artificial Intelligence **17** (1981) 409-460.

24. Brady, M. and Yuille, A., An extremum principle for shape from contour, IEEE Trans. Pattern Anal. Mach. **6** (1984) 288-301.

25. Liang, P. and Todhunter, J.S., Three dimensional shape reconstruction from image by minimum energy principle, Proceedings 2nd Conference on Artificial Intelligence Applications (1985) 100-105.

26. Yuen, S.Y., Computer perception of hidden lines, MPhil thesis, Dept. of Electronic Engg., Hong Kong Polytechnic (1988).

27. Marr, D., Early processing of visual information, Phil. Trans. R. Soc. Lond. **275**(942) (1976) 483-534.

28. Attneave, F., Prägnanz and soap bubble systems : a theoretical exploration, in: J. Beck (Ed.), Organization and Representation in Perception (Lawrence Erlbaum Associates, 1982) 11-29.

29. Terzopoulos, D., Witkin, A. and Kass, M., Constraints on deformable models : recovering 3D shape and nonrigid motion, Artificial Intelligence **36** (1988) 91-123.

30. Hogg, D.C., Interpreting images of a known moving object, DPhil thesis, Univ. of Sussex (1984).

31. Friedberg, S.A., Finding axes of skewed symmetry, Comput. Vision Graph. Image Process. **34** (1986) 138-155.

32. Ponce, J., Ribbons, symmetries and skewed symmetries, Proceedings Image Understanding Workshop (1988) 1074-1079.

33. Scott, G.L., Turner, S.C. and Zisserman, A., Using a mixed wave/diffusion process to elicit the symmetry set, Proceedings of the 4^{th} Alvey Vision Conference (1988) 221-228.

34. Mumford, D., The problem of robust shape descriptors, Proceedings 1^{st} International Conference on Computer Vision, London (1987) 602-606.