same for all tested images. However, the robustness of the method to changes in those parameters has been tested by changing their values up to 50% of their default ones. The changes have not affected the final results although they produce larger search spaces for looser values (table Table 1: given in section 4.4.3 contains the numbers of hypotheses at each step using the default setting of the parameters). The fact that the same final results are obtained is due to the hypothesize-verify nature of the segmentation method whereby final decisions are made based on global evidence of regularity.

4.7 Conclusion

This work shows that the problem of segmentation and recovery of 3-D curved objects from a single real image can be solved for a large class of objects. It shows that there is great advantage in viewing image descriptions as *projections of 3-D descriptions* rather than as just 2-dimensional ones. This view leads to the analysis and usage of projective properties of 3-D shapes for their segmentation and recovery. The geometric quasi-invariant properties provide means to estimate the projection of 3-D descriptions which can then be used to recover 3-D shape. The structural properties allow to both predict arrangements of features in the image, which enables the method to handle feature fragmentation, and generate consistency constraints in hypothesizing (3-D) objects from a clutter of image contours.

The discussion illustrates the issues and difficulties in addressing the 3-D shape segmentation problem from a real intensity image. It also gives methods that can be used in solving them. However, we have only addressed some curved axis primitives in this paper. In a previous effort [81], we have derived a similar method for segmentation and recovery of SHGCs. These two methods, together, are the core of a larger system for segmentation and recovery of compound objects made up of several components of both types (a large class of objects). Such compound objects introduce difficulties that have not been addressed in this work. Due to joints and the complex arrangement of components, geometric as well as structural regularities may not be observed in the image. We are currently investigating generic solutions to this problem.

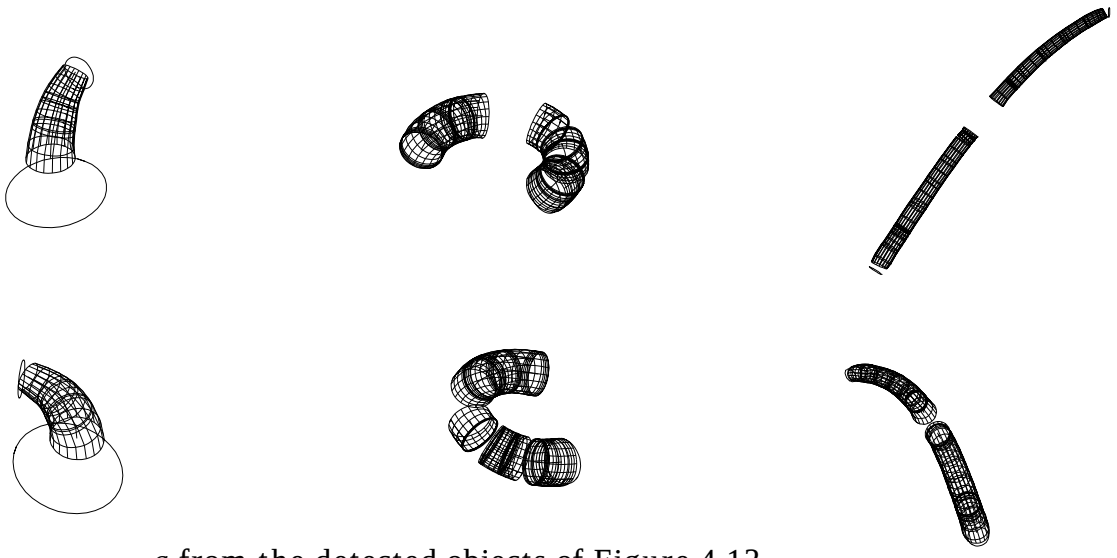

*a.*from the detected objects of Figure 4.13

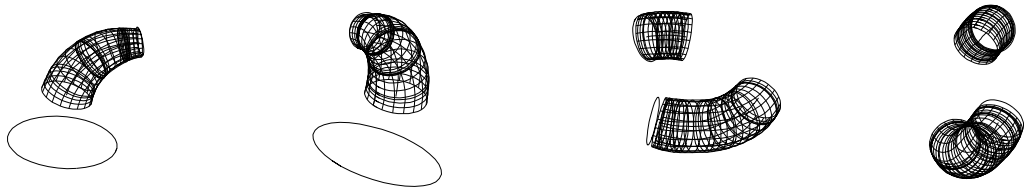

b. from the detected objects of Figure 4.14

c. from the detected objects of Figure 4.14 (continued).

Figure 4.26 Recovered 3-D primitives shown for their original poses and different ones.

patterns discussed in section . The given examples show instances of self-occlusion closure (the left-most one in the images) and the *3-tgt-j 3-tgt-j* closure which is also not handled in [62].

The method of [45] uses an intuitive definition of a ribbon which necessarily matches the extremities of a pair of curves. Here, we have used rigorous projective quasi-invariants and invariants for justifying the choice of right ribbons for finding the correspondences in the image.

The method we have described uses a number of parameters to account for noise, quantization and image feature discontinuities. The parameter values have been the

consists, first, of estimating the 3-D orientations of the circles that project onto the detected elliptical cross-sections (that can be partially visible). Since the cross-section has been classified as pointing either towards or away from the viewer, there a unique solution for each ellipse. The 3-D axis plane orientation is thus the cross-product of the obtained circle orientations. Then, at each location of the visible surface the recovery consists of finding the orientation and scale of the 3-D circle whose projection passes through that location. Details of the method can be found in [82].

However, it was applied to perfect and segmented contours. Here, we demonstrate its application on automatically detected objects from real images. The 3-D recovery is applied to each visible surface patch of a detected circular PRGC. Discontinuity regions between surface patches cannot be uniquely recovered (unlike the case of SHGCs where there is a unique way [81]).

Results of the application of the 3-D recovery method to the segmented objects of Figure 4.24 are shown in Figure 4.26.*a* with the ruled 3-D surfaces in terms of meridians and cross-sections shown for different 3-D orientations. Notice, for the first object, the gap between the bottom (larger) cross-section and the surface which is due to both self-occlusion and low contrast in the image which prevented visibility of the terminating part of the limb boundaries of the object. Also the other objects have discontinuities in their recovered 3-D descriptions due the (self) occlusions and low contrast discontinuities in the image.

Figure 4.26.*b* and Figure 4.26.*c* show the resulting 3-D descriptions obtained from the detected objects of Figure 4.25 and Figure 4.25 (continued) respectively. Here too the discontinuities in the 3-D descriptions are due to the self-occlusion (left object in case *b.* and both objects in case *c.*), occlusion (right object in case *b.*) and markings-caused discontinuities in the image (left object in case *c.*). Notice that the recovered 3-D horn-like object (left-most in all images) is the same although it has different poses in the images.

4.6 Discussion

The method has been fully implemented and tested on several images. We have given three of them in this paper. The examples shown illustrate the difficulties with real image imperfections such as surface markings, shadows, contour breaks and occlusion. Methods based on simple contour following (such as [3] for e.g.) would fail to handle those imperfections. As we have previously mentioned, previous ribbon-based methods do not rigorously address the complexities of the projection of curved objects as they use intuitive methods and solve a 2-dimensional problem. For example, closure verification in [62] uses a simple sequence of curves between the extremities of a ribbon with the constraint that they remain “outside” the ribbon and within its width. Although, this handles the *L-j L-j* pattern, it does not handle most of the other patterns. Our method explicitly accounts for the structural complexities of the projection of 3-D curved objects. In the case of closure, for example, it searches for all the

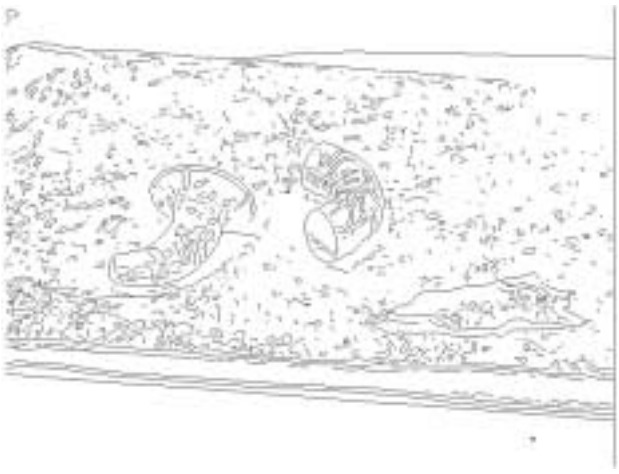

a. initial intensity image.

b.edge image

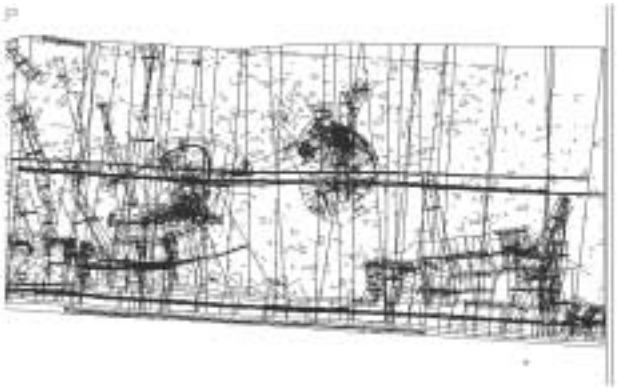

c.some detected local surface patches

d. verified objects

Figure 4.25 (continued)Additional example of results of the segmentation method.

Table 1: summary of results for the given examples

image	numb. local. surface patches	numb. object. hypotheses.	numb. verified objects	numb. verified local patches
Figure 4.1	90	69	3	8
Figure 4.25	167	63	2	3
Figure 4.25 (continued)	298	197	2	4

4.5 3-D shape recovery

In our previous work [82] we have shown how to exploit the geometric quasi-invariant properties of circular PRGCs to recover their 3-D object centered descriptions in terms of the 3-D cross-section, 3-D axis and the 3-D scaling function. The method

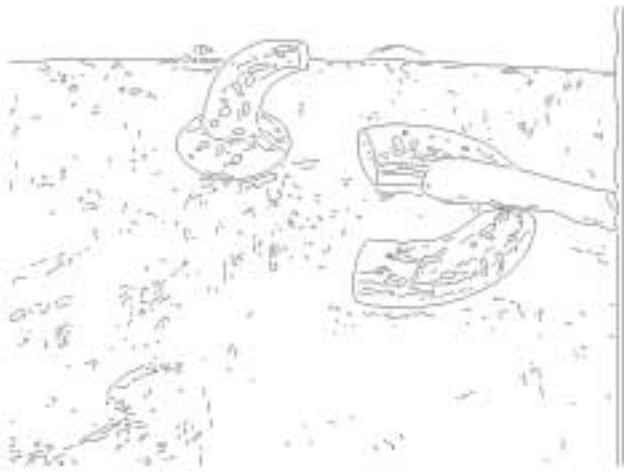

a.initial intensity image

b. edge image

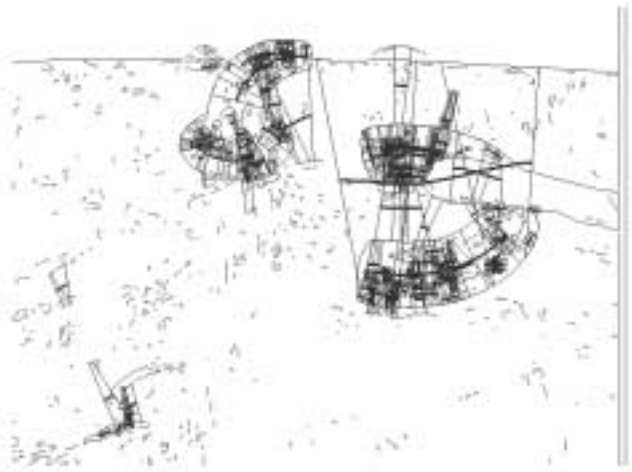

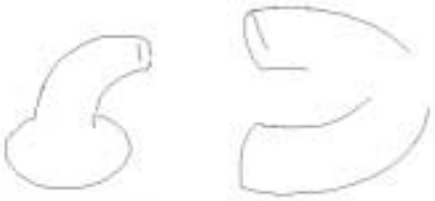

c. detected local surface patches

d.verified objects.

Figure 4.25 Additional example of resulting object hypotheses for a different scene of similar objects.

objects with the corresponding number of the verified local surface patches of the original set (i.e. the component local surface patches of the verified objects)

Figure 4.24 Verified object hypotheses from the surface patches of Figure 4.14.

Figure 4.24 shows the verified object hypotheses from the patches of Figure 4.14. The left-most object has been identified as having an *L-j L-j* closure at its top part and a self-occlusion closure at its bottom part (the pattern of Figure 4.8.*c*) and both cross-sections are classified as pointing away from the image. The middle object has been identified as having *3-tgt-j 3-tgt-j* closures at both extremities and both cross-sections have been classified as facing the image. The right cross-section has initially been classified *3-tgt L-j* but the *L-j* mislabelling has been corrected by finding the full cross-section closure. The right-most object has been identified has having *L-j L-j* closures at both extremities and both cross-sections have been classified as pointing away from the image.

Figure 4.25 and Figure 4.25 (continued) show additional results of the method applied to different scenes of similar objects but with more markings interacting with their surfaces and different poses. The piece of paper occluding the right-most object in Figure 4.25 has not been labeled as a circular PRGC because its contours do not produce parallel symmetry (the upper part is straight and the lower part has many irregularities; its shape has a straight axis and straight axis primitives are not addressed here but in [81]).

Table Table 1: summarizes the results of the evaluations for the images of Figure 4.1 (and its results of Figure 4.14 and Figure 4.24) and of Figure 4.25 and Figure 4.25 (continued). The columns give the number of hypothesized local surface patches, the number of object hypotheses (after grouping) and the number of verified

a. original boundaries *b.* L-junction L-junction closure *c.* T-junction T-junction closure

Figure 4.22 Several junctions may be found at a hypothesized object's end

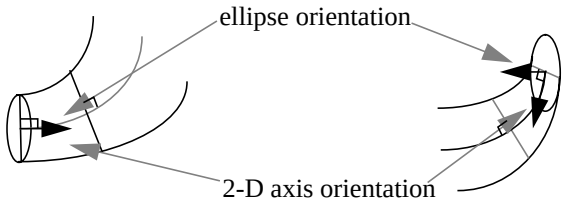

a.coincident in a regular case *b.non-coincident in an irregular case*

Figure 4.23 Extremal cross-section and axis orientations are constrained in the projection of circular PRGCs.

onality in 3-D between the cross-section and the axis of a circular PRGC and by virtue of properties QP2 and IP2, the ellipse minor axis should be (almost) parallel to the ribbon axis at that extremity (the circle and axis orientations are the same in 3-D, the image ellipse minor axis is the projection of the normal to the circle plane and the right ribbon axis is close to be parallel to the projection of the axis tangent). This test is only applied to complete descriptions so as not to reject incomplete (but right) ones. Closed object hypotheses that do not have elliptic cross-sections or that are complete but do not satisfy the above test are marked as "unknown" although they are considered as meaningful scene objects.

The closure patterns found for an object hypothesis can be used to classify the cross-section (when visible) as either pointing away or towards the image plane. The cross-section is classified to be pointing towards the image plane if any of the closure patterns of Figure 4.7 is found, namely either a *3-tgt-j* is found or two *T-j*s with evidence of a (possibly partially occluded) ellipse between them. It is classified to be pointing away from the image plane otherwise.

Notice finally that the *3-tgt-j L-j* combination, which is not possible in theory (as seen in section), might be detected in a real image due to low contrast in the vicinity of the cross-section which causes the detection of an *L-j* instead of a *3-tgt-j*. In this case, the method reconsiders the found *L-j* and searches for evidence of a visible cross-section in the same way as for a *3-tgt-j 3-tgt-j* pairing.

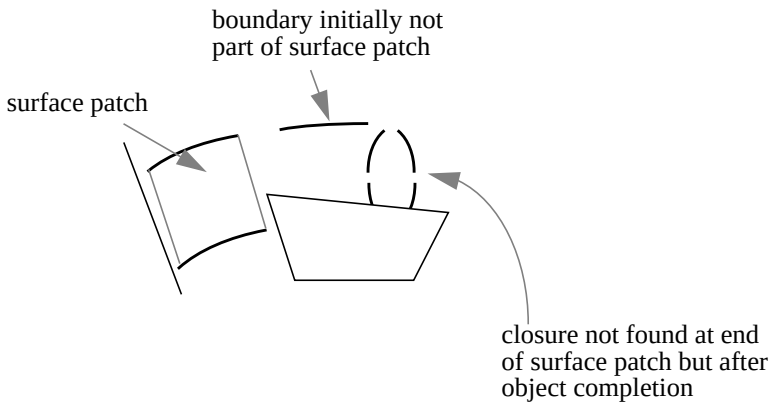

Figure 4.20 Closure finding and completion of object description

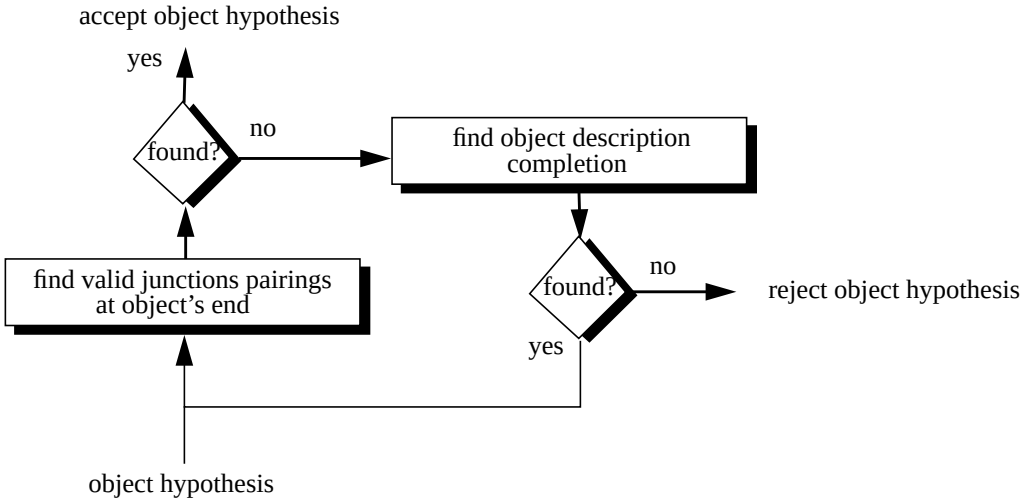

Figure 4.21 Block diagram of the closure-completion test

Due to surface markings, shadows and the use of junction measures, several junctions may be found at an end of an object hypothesis. For example, in Figure 4.22, two types of closures namely *L-j L-j* and *T-j T-j* are detectable. In this case, the closest of the junction pairings found is selected.

The search for junctions and connectivity between boundaries is efficiently implemented in our system by using spatial indexing of image boundaries whereby, given a desired region the image, the boundaries belonging to it are directly accessed (usually a small number; an average of 6 boundaries per 32 square pixels for the image of Figure 4.25).

In case the detected object cross-section is visible, it is tested whether it is elliptical or not (using fitting criteria). If it is found to be elliptical, then from the orthog-

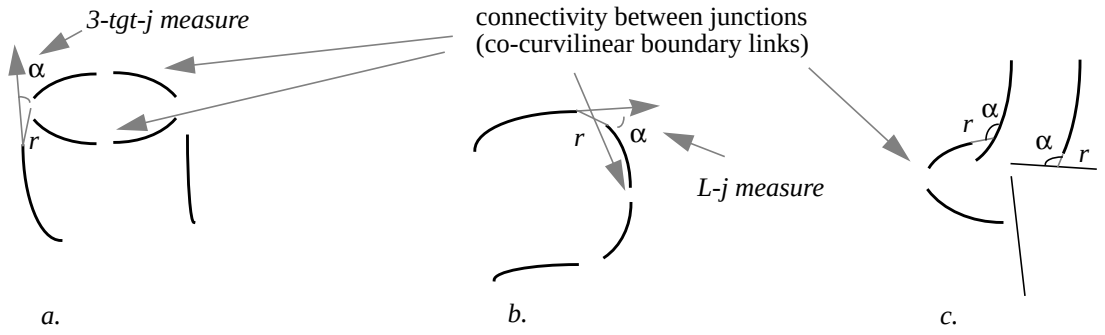

Figure 4.18 Closure finding: junction measures and boundary links.

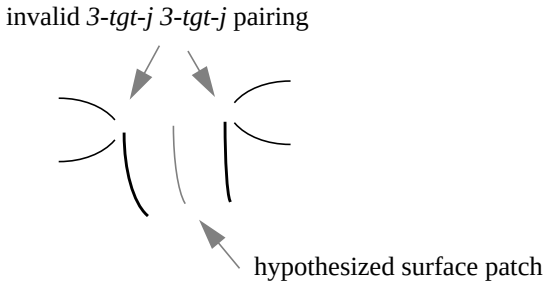

Figure 4.19 example of invalid closure junction pairings resulting in a non-closed object hypothesis

proximity and angular measures (see Figure 4.18.*a* and *b*). Valid junction pairings are those for which the connectivity relationships of section 2.2.2, depending on the type of junctions, are found. Connectivity relationships between the junctions use co-curvilinearity measures between boundaries but with looser thresholds (allowing larger gaps and angular variations) than those of the curve level to account for its false negatives. Except for the case of a completely occluded cross-section, the sequence of boundaries between the junctions is hypothesized to be the cross-section projection (Figure 4.18).

Object hypotheses for which none of the closure patterns of section is found are rejected. An example is given in Figure 4.19. Here, the (wrong) object hypothesis has a *3-tgt-j 3-tgt-j* pairing but no connectivity is found between the junctions and the hypothesis is thus ruled out.

Since the curve level grouping may have false negatives, it is possible that no junctions are detectable at a hypothesized object's end (Figure 4.20). In this case, the method attempts to complete the object's description by finding boundary links (also with looser thresholds) that produce a valid object closure. In doing so, the system is more robust to false negatives of the curve level grouping, a desirable property of a segmentation method. Figure 4.21 shows the block diagram of the search for closure and object completion process.

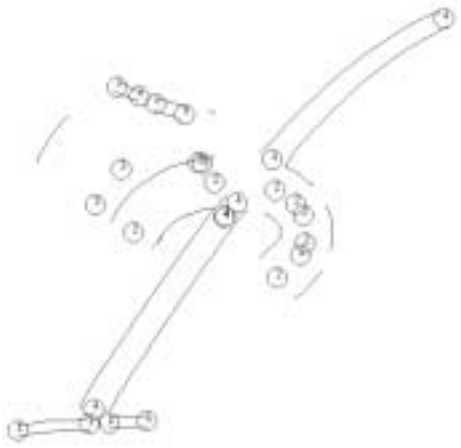

Figure 4.17 Examples of grouping hypotheses from the surface patches of Figure 4.14.c.

smooth functions, the grouping hypothesis selection consists of choosing the surface patch which minimizes the following measure:

$$S_I = r_I \, e_I \, m'_I$$

where r_I is the size of the relative gap between the surface patches, e_I a measure of relative axis curvature between them (the projective counterpart of the measure e of section) and m'_I a measure of the relative change of the cross-section segment size (the projective counterpart of the 3-D scaling derivative $\dot{r}$ of section ; itself constrained to be small enough). This measure maximizes shape smoothness. Figure 4.17 shows some of the grouping hypotheses for the surface patches detected in Figure 4.14. Patches of the same group are labeled with the same number.

4.4.3 Verification of object hypotheses

The previous steps do not guarantee that a hypothesized object (aggregate of surface patches of possibly one element) is the projection of a meaningful object in the scene because the properties we use are necessary projective properties but not sufficient ones. The surface patches shown in Figure 4.14 include many examples of wrong hypotheses.

The verification method consists of imposing closure at each extremity of an object hypothesis. For this, the method searches for any of the closure patterns discussed in section which consist of junction pairings and connectivity relationships between them (i.e. the cross-section projection or evidence of its occlusion). To find those closure patterns, and since no prior knowledge of the object or its pose is given, the method first searches for evidence of *3-tgt-j*s, *L-j*s and *T-j*s involving the hypothesized object's boundaries in the vicinity of each of its extremities. To account for real image imperfections, junction and connectivity measures are used. As for the *T*-junction measure previously shown in Figure 4.15.*a*, detection of *3-tgt-js* and *L-js* also use

Figure 4.16 constraints on discontinuous surface patch grouping

junction among each other). Their identification consists of constraining the relative space, *r,* between the boundaries and the variation of their end tangents to be small enough. The smoothness measure used, *M,* is as shown in Figure 4.15.*b*. This constraint corresponds to the expectation of both proximity and orientation continuity of the surface in the image.

Two local surface patches are considered to be *compatible* under one of the following cases (using the terminology of the relationship between patches given in section):

- they are *continuous*
- they are *self-occluding* with the constraint that the occluded boundaries belong to the *same surface patch* and the cusps are *internal* to the other surface patch
- they are *discontinuous* with the constraint that a conic exists between both pairs of matched boundaries (in addition to the co-curvilinearity constraint mentioned above; Figure 4.16.*a*)

In the first case above, boundary continuity between the surface patches is an indicator of surface continuity. In the second case, the constraint on the occluded boundaries follows from properties ISP2 and ISP3 of section and avoid the inconsistencies of Figure 4.6. In the third case, imposing a conic to exist between the extremities (with co-curvilinearity of the boundaries) avoids grouping the patches of Figure 4.16.*b* for which no conic exists for the top boundaries (since an inflection would be necessary). In other words, we are assuming that the surface is not fragmented (or occluded) at its inflection regions. Surface patches that violate these constraints are not considered for grouping.

Ribbon continuity has also been used in the methods of [34,62] for grouping ribbons. However, in those methods, compatibility constraints do not take account of the self-occlusion complexities we have described; *i.e.* their visibility model does not include complex arrangements of boundaries. The constraints we have described make explicit the fact that projections of 3-D objects are at hand.

Finally, for a local surface patch, at each of its ends, more than one compatible local surface patch may be found. Since the 3-D axis and sweep of a circular PRGC are

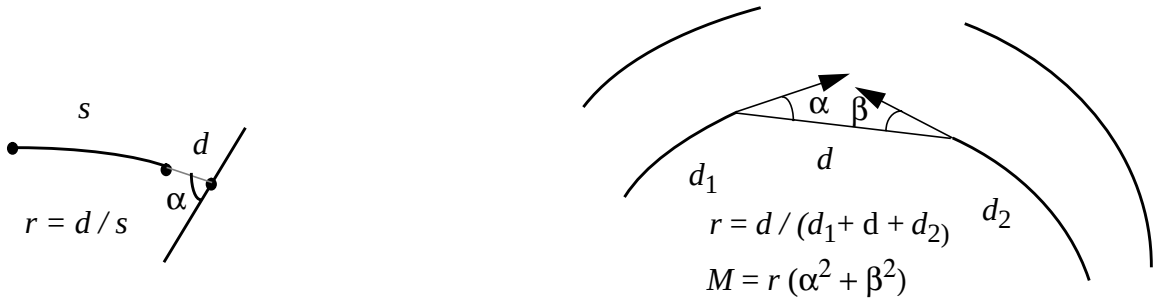

a. T-junction measure

b. smoothness measure for discontinuous patches

Figure 4.15 Finding structural relationships between boundaries of local surface patches

8 are fragments of meaningful circular PRGCs). This is because the geometric projective properties of section 4.2.1 are *necessary* conditions of the projection of circular PRGCs but *not sufficient* ones to firmly conclude their presence. Furthermore, due to the sources of image discontinuities discussed in section 4.2.2, several local surface patches may be produced for the same scene object. Grouping of local surface patches belonging to the same object and eliminating the wrong ones are discussed in the next two steps.

4.4.2 Grouping of local surface patches

In section , we have discussed the different relationships between surface patches of the same object. The contours of two such surface patches are either the same, form self-occlusion discontinuity or occlusion discontinuity. In a real image, discontinuity also occurs due to errors in edge localization or contour breaks at low contrast regions as is the case of the occluding tube of Figure 4.1. We will call this type of discontinuity a *contour-break discontinuity*. To analyze the relationship between two local surface patches detected in the previous step, we have to determine whether their contours form any of the structural patterns of section , including contour-break discontinuity.

Identifying the continuity relationship is trivial since it corresponds to having the same boundary. Identifying self-occlusion discontinuity between a pair of different boundaries consists of finding evidence of a *T*-junction between them. *T*-junctions are not expected to be perfect; i.e. there may not be actual contact between the boundaries due to, say, low contrast. For this, a *T*-junction measure is used and consists of the relative distance, *r*, between the boundaries (ratio of the length of the gap to the length of the stem curve of the *T*) and the angle, α, of contact (Figure 4.15.*a*). The former should be small and the latter large enough to rule out collinear arrangements. In the implementation, the thresholds used are 0.2 for *r* and 15^o for α.

Identification of occlusion discontinuity and contour-break discontinuity between the boundaries is the same (the boundaries are different and do not form a *T*-

a. finding extremal correspondences *b. finding subsequent correspondences*

Figure 4.13 Right ribbon detection.

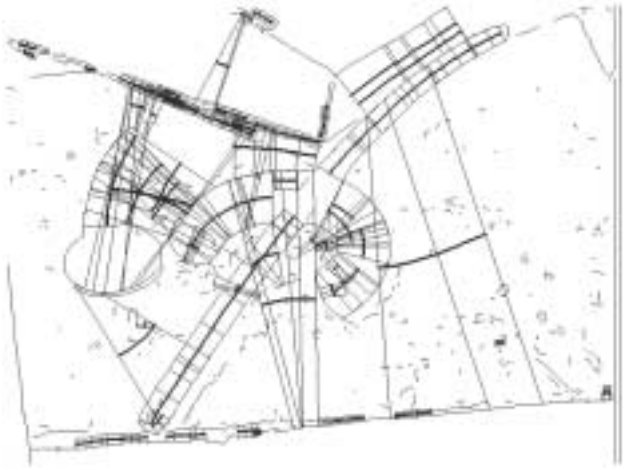

Figure 4.14 Resulting local surface patches from the image of Figure 4.1.

extension method uses the same minimization method for finding right ribbon correspondences.

Since parallel symmetry already gives the co-cross-section limb points of the projection of circular PRCGCs (constant sweep), it is not necessary to use the above method in such cases. To differentiate between the constant cross-section case from the non-constant cross-section case, we can use invariant property IP3 which indicates that the size of the cross-section segments is constant in the projection of a circular PRCGC. Thus, for parallel symmetries with (almost) constant length correspondence segments, the search for right ribbons is not performed and only an orthogonality test between correspondence segments and axis tangents is needed. We will call *local surface patch* the description given by a pair of image contours and their right ribbon correspondences so determined.

Figure 4.14 shows some of the local surface patches so detected from the contours of Figure 4.1. The figure displays cross-section segments and axes (in bold lines). It is clear that the method does not yield only local surface patches that project from actual scene objects (the total number of hypothesized surface patches is 90, only

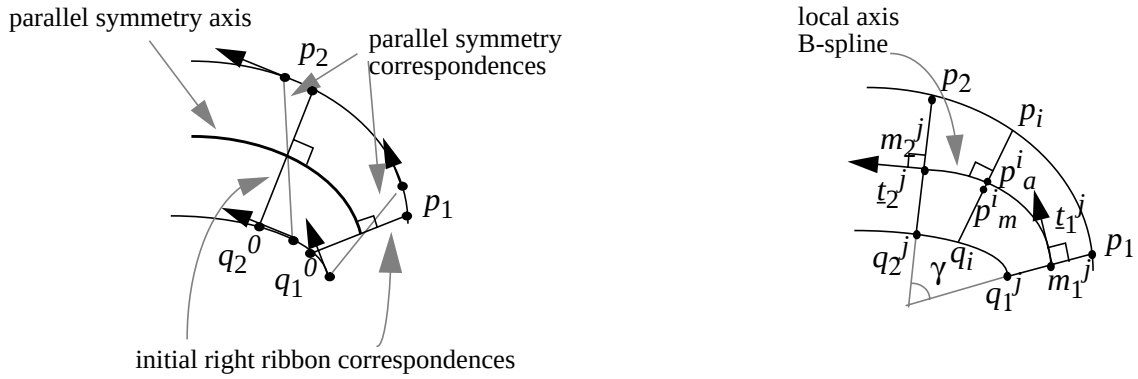

a. *b.*

Figure 4.12 Local right ribbon detection. a. parallel symmetry axis as initial estimate of right ribbon axis. b. finding the local axis B-spline.

s_2^j (see Figure 4.12.*b*). The "best" such right ribbon segment is defined as the one which minimizes the following measure:

$$E = \lambda_1 E_1 + \lambda_2 E_2 \tag{4.1}$$

where $E_1 = \dfrac{\sum_i^n dist(p^i_a, p^i_m) / dist(p_i, q_i)}{n}$; p^i_a is an i^{th} point on the local axis, p_i, q_i, corresponding points (intersections of the line orthogonal to the axis at p^i_a with the two curves) and p^i_m the midpoint of $p_i\, q_i$ (Figure 4.12.*b*)

and $E_2 = \gamma / dist(m_1^j, m_2^j)$; γ being the angle between the segments s_1^j and s_2^j.

Minimizing E_1 forces the distance between the local axis and the mid-points of corresponding segments $p_i\, q_i$ to be small; i.e. the local axis should be the 2-D axis as defined in section 4.2.1. Minimizing E_2 avoids high curvature axes which would correspond to highly curved and thick shapes that are close to self-intersect [82]. λ_1 and λ_2 are real scalars (whose sum is 1). The method gives more weight to E_1 than E_2 as the former *should* be zero for a right ribbon axis and the latter is only a regularity measure (in our system λ_1 is fixed to 0.8 and λ_2 to 0.2).

For each pair of parallel symmetric contours, this search is initially carried out for the first two points at an extremity of the parallel symmetry axis at hand (p_1 and p_2 in Figure 4.13.*a*). The resulting correspondences (if any) are $p_1\, q_1$ and $p_2\, q_2$. Then the search is performed at one point at a time (p_k in Figure 4.13.*b*), along one of the contours, using the previously found right ribbon segment ($p_{k-1}\, q_{k-1}$) to minimize the measure E of equation (4.1). The search region is small as the parallel symmetry axis is in general close to the right ribbon axis.

This process if carried out up to the extremities of the parallel symmetry axis. An *extension* of the obtained right ribbon is performed so as to match at least one curve extremity (the parallel symmetry axis may not match any curve extremity). The

discussed there, includes only the methods for forming SHGC descriptions. In section 4.4, we will discuss the surface patch level of the part of our system that handles curved axis GCs. Although this paper addresses mainly circular PRGCs, most of the methods presented apply to other curved axis primitives as well.

4.4 Detection of circular PRGCs (surface patch level)

The detection of circular PRGCs from real image contours consists of a multi-step hypothesize-verify process: the detection of local surface patches, the grouping of those projecting from the same object and the verification of object hypotheses. 3-D shape is then recovered for verified object hypotheses whenever possible. In this section, we discuss the details of the three steps of the segmentation process. The 3-D recovery is discussed in section 4.5.

4.4.1 Detection of local surface patches

From the discussion in section 4.2.1, right ribbons give a good approximation of the projection of the 3-D description of a circular PRGC and an exact one for the circular PRCGC sub-case. Thus, a first step to recovering a circular PRGC is the detection of right ribbons. In [50] a right ribbon detection method (*projection method*) has been proposed and consists of discretizing the orientations of the axis and finding correspondences for each orientation. The complexity of this method is $O(km)$ where k is the number of discretized axis orientations and m the number of image points, which is a costly method for highly cluttered scenes. A possible improvement would use the B-spline representation of image contours and find local ribbon correspondences between B-spline segments [82]. The complexity is also $O(km)$, where m is now the number of B-spline segments. However, in a real image that contains a large number of contours that method is still costly.

From the discussion in section 4.2.1, we have also seen that parallel symmetry correspondences exactly coincide with right ribbon correspondences for the constant sweep case and that in the case of a non-constant sweep there is an offset between them. Since parallel symmetries are efficiently detected in our system, our right ribbon detection method between a pair of curves consists of using their parallel symmetry axis (if any) as an initial estimate of the right ribbon axis in a minimization method. The method consists of searching around the initial correspondences ($p_1q_1{}^0$ and $p_2q_2{}^0$ in Figure 4.12.*a*) for the "best" right ribbon segments between two successive quadratic B-spline extremities (p_1 and p_2). The search is over an angular region $\Delta\theta$ centered at the initial correspondences (Figure 4.13). At each step j of the search, the correspondences are the segments $s_1{}^j = p_1\ q_1{}^j$ and $s_2{}^j = p_2\ q_2{}^j$. The *local axis* is defined to be the quadratic B-spline segment (if it exists) defined by $m_1{}^j$ and $m_2{}^j$, the midpoints of $s_1{}^j$ and $s_2{}^j$, and orientations $\underline{t}_1{}^j$ and $\underline{t}_2{}^j$ which are the normals to $s_1{}^j$ and

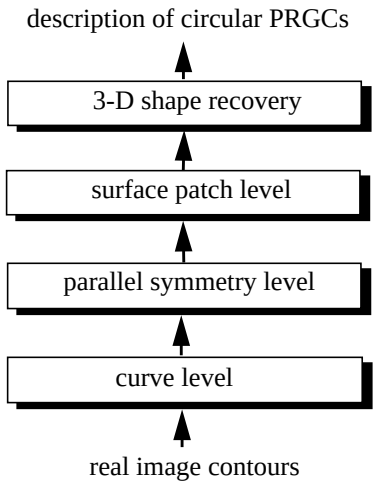

Figure 4.10 Block diagram of the segmentation method

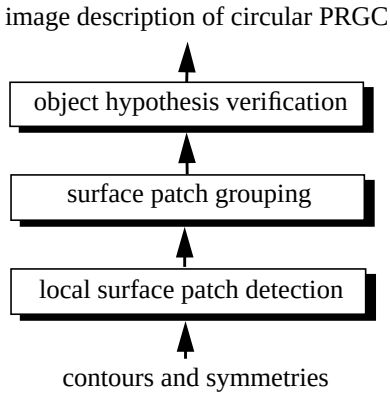

Figure 4.11 Block diagram of the surface patch level

The symmetry level is intended to detect parallel symmetry relationships between contours produced by the curve level. Detection of parallel symmetries uses a quadratic B-spline representation of image boundaries [66] and is thus analytic and efficient. It results in piecewise continuous correspondences between pairs of curves. Those symmetries allow to form initial hypotheses about the presence of scene objects (GCs) that are either confirmed or ruled out at higher levels.

The surface patch level is intended to form complete GC descriptions whenever possible in the image. It consists of searching for evidence of surface fragments, grouping compatible ones together and verifying them on the basis of expected projective properties derived from the geometric invariants and quasi-invariants and structural regularities such as object closure. Verified object hypotheses are then used to recover 3-D shape if there is sufficient information in their image description. Figure 4.11 shows the block diagram of the surface patch level.

The curve level and the symmetry level are discussed in detail in a description of the part of our larger system that handles SHGCs in [81]. The surface patch level,

a. convergent closure *b. T-j T-j closure*

Figure 4.9 Image closure patterns when the cross-section is not visible.

3) Non visible cross-section

There are two cases.

- the cross-section scaling vanishes: then the two surface boundaries meet at a point (*convergent* closure; Figure 4.9.*a*)
- there is occlusion by another surface: then two *T-j*s (where the surface boundaries being the *stems* of the *T*'s) are visible in the image (Figure 4.9.*b*).

The case of joints between objects is also a source of non-visibility of the cross-section. In this case surface closure consists of the intersection curve of the surfaces of the joined objects. Compound objects are not considered part of our analysis.

The above analysis, besides giving closure tests for automatic object verification, also allows inference of useful information about the shape of the hypothesized object at hand. For example, classifying the cross-section as facing the image plane or pointing away from it, using the observed closure pattern, helps finding its correct 3-D orientation and therefore the whole object's 3-D shape and pose. Usage of this analysis and the effects of contour breaks are discussed in section 4.4.3.

In traditional ribbon-based segmentation methods which use surface closure for verifying objects hypotheses, simple closure criteria were used, typically the existence of a single path of curves joining the extremities of a ribbon [62]. However, such criteria do not capture the more complex closure patterns and junctions we have described. The inner surface and surface closure properties we described explicitly take into account that the objects sought are 3-dimensional.

4.3 Overview of the approach

Our approach to the figure-ground problem for GCs consists of a hierarchy of feature levels: curves, symmetries and surface patches (Figure 4.10). The curve level consists of forming contours from edges obtained by an edge detector (such as [14] for example). Local contour fragments are first formed by simple edge linking operations, then grouped into global contours using a conservative co-curvilinearity based grouping aimed mainly at handling short breaks such as usually occur due to poor edge localization.

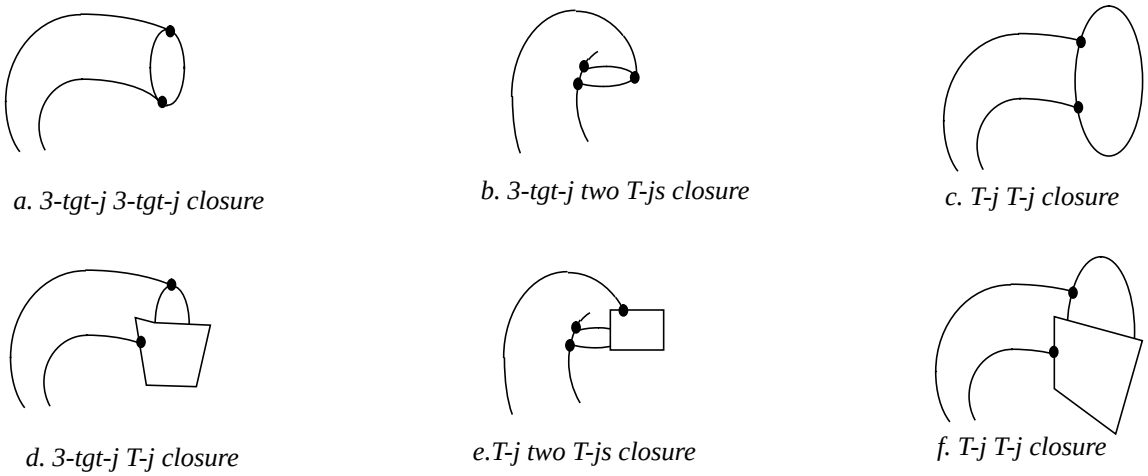

Figure 4.7 Image closure patterns when the cross-section faces the viewer.

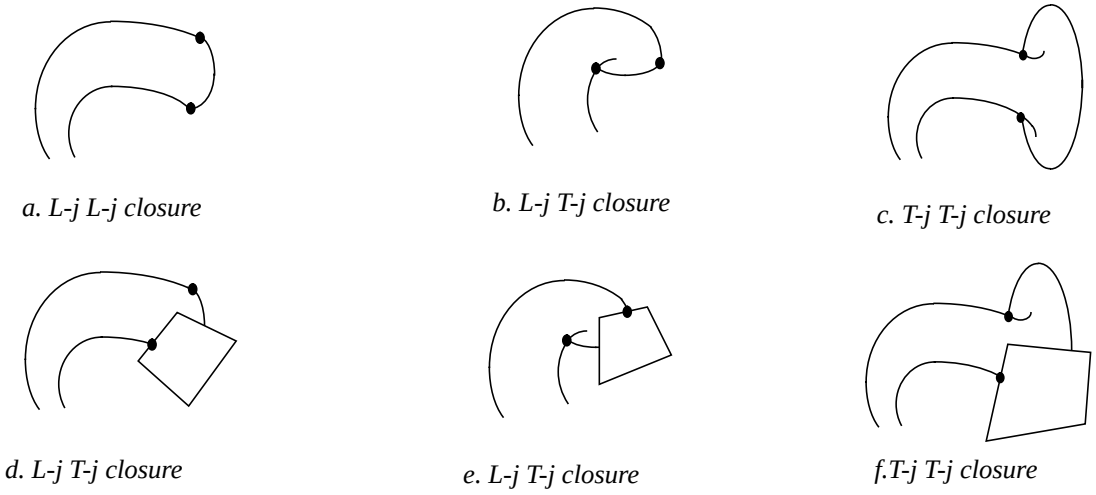

Figure 4.8 Image closure patterns when the cross-section points away from the viewer.

The three combinations are *L-j L-j*, *L-j T-j*, *T-j T-j*. The sources of those combinations can be classified into three cases:

- there is no occlusion at all: then two *L-j*'s are visible in the image with an elliptic arc joining them (Figure 4.8.*a*).
- there is only self-occlusion: in this case, surface cusps are visible in the image and two patterns are possible: one *L-j* and one *T-j* with an elliptic arc between them (Figure 4.8.*b*) and two *T-j*'s (where the surface boundaries being the *top*s of the *T*'s) with an elliptic arc joining them (Figure 4.8.*c*).
- there is occlusion by another surface: then at least one *T-j* (where the surface boundary is the *stem* of the *T)* is visible with partially visible elliptic arcs (Figure 4.8.*d-f)*.

Surface closure properties

In 3-D, a GC is terminated by its cut. In the image, depending on the viewing direction, the object's shape and other scene and image complexities, several junctions terminate its projection. The terminology of the junctions we use is borrowed from [40]. Our purpose in this paper is *not* to use junction labelling as the only means to infer 3-D structure (in the sense of [40] for example). Rather, we want to predict the different termination patterns of the projection of 3-D objects. The intent, as we will discuss in section 4.4.3, is to derive *verification* tests for object hypotheses.

Three types of junctions may be observed in the projection of a circular PRGC: *three-tangent* junction (or simply *3-tgt-j*) which corresponds to a cross-section facing the image plane, *curvature-L* junction (or simply *L-j*) which corresponds to a cross-section pointing away from the image plane and *T*-junction (or simply *T-j*) which results from an occlusion (including self-occlusions with cusp endings [37]). The closure patterns and the associated junctions at the end of the boundaries of a surface depend on whether or not its cross-section is visible and in case it is, whether it points towards or away from the image plane (viewer). For each of the latter two cases, there are three combinations of the above junctions since the *L-j* is not observed in the first case and the *3-tgt-j* is not observed in the second. In what follows, we give the closure structural properties as a list of cases, which we believe to be complete, consisting of junction combinations and connectivity relationships between them (see Figure 4.7 to Figure 4.9).

1) Visible cross-section pointing towards the viewer

The three combinations are *3-tgt-j 3-tgt-j*, *3-tgt-j T-j* and *T-j T-j*. The sources of those combinations can be classified into four cases:

- there is no occlusion at all: then two *3-tgt-j*'s are visible in the image with an ellipse joining them (Figure 4.7.*a*).
- there is only self-occlusion with the surface partially occluding the cross-section: then a surface cusp is visible and we obtain the pattern of Figure 4.7.*b* which consists of one *3-tgt-j* and two *T-j*s with the same boundary (where the surface boundary is the *top* of the *T*) and a partially occluded ellipse joining the two sides.
- there is only self-occlusion with the cross-section partially occluding the surface: then we obtain the pattern of Figure 4.7.*c* which consists of two *T-j*s (where the surface boundaries being the *stems* of the *T*'s) with an ellipse joining them.
- there is occlusion by another surface: then at least one *T-j (*where the surface boundary is the *stem* of the *T)* is visible with a partially occluded ellipse (Figure 4.7.*d-f).*

2) Visible cross-section pointing away from the viewer

Figure 4.6 impossible structural arrangements of self-occluding surface patches.

The structural relationships between the surface patches (fragments) of the projection of an object can be expressed as combinations of the relationships between their contours discussed above. Considering an image surface as bounded by two boundaries, there are six combinations of the above contour relationships (Figure 4.5.*a* through *f*). We will call such surface patches *continuous* if they share a contour and an occlusion discontinuity holds among the other two (Figure 4.5.*b*). We will call them *self-occluding* if a self-occlusion discontinuity has been found for their contours (Figure 4.5.*c*, *e* and *f*) and we will call them *discontinuous* if only occlusion discontinuities have been found for their contours (Figure 4.5.*d).* Thus the inner surface property below holds.

Property ISP1: A surface either projects onto a single surface patch or onto many surface patches related by the above structural events (self-occlusion discontinuity or occlusion discontinuity).

There are also straightforward (but powerful) properties of the combinations involving self-occlusion discontinuity. They are given below.

Property ISP2: In the image, a surface cusp is always *inside* the region between the boundaries of the projection of an object.

Thus the pattern of Figure 4.6.*a* is impossible.

Property ISP3: In case both image boundary relationships are self-occlusions, the cusps belong to the same surface patch.

Thus the pattern of Figure 4.6.*b* is also impossible.

The above analysis is useful in that it allows to derive surface patch grouping constraints in a method for automatic detection of circular PRGCs projections in an image. The last simple consistency constraints are useful in that they provide filters to rule out inconsistent groupings. This is discussed in section 4.4.2.

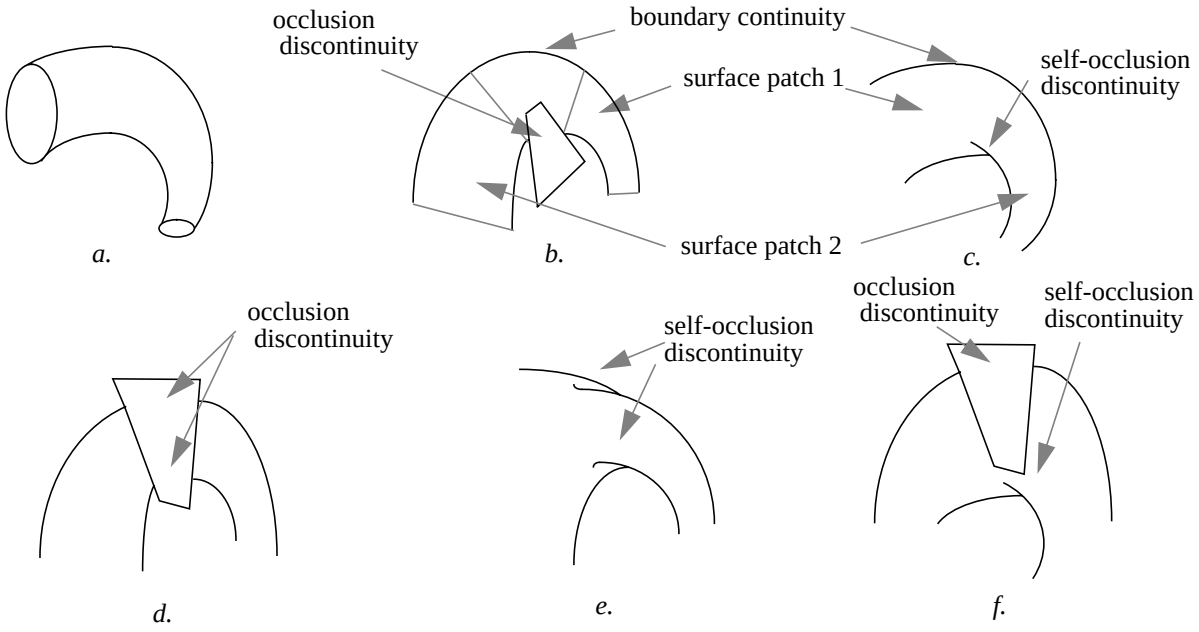

Figure 4.5 Inner surface structural properties

We distinguish two types of structural properties: those that give relationships between contours of the swept surface of the same object (called *inner surface properties*) and those that hold at an object's end (called *surface closure properties*). For lack of space, the mathematical formulation of the shape measures and viewing angles for which the structural events of circular PRGCs occur will be omitted (it is similar to the one given in [71] for SHGCs) and we give a qualitative description of those events (see also [56] for derivation of limbs and cusps of GCs and [37] for an analysis of *generic events* of curved shapes). Effects of real image imperfections such as contour breaks will be discussed in section 4.4 where we describe the segmentation method. The properties assume a non degenerate viewing direction.

Inner surface properties

Due to self-occlusion, an object's surface may project onto several *surface patches*. The relationships between the boundaries of those surface patches can be classified into three categories (see Figure 4.5): the *continuous* case, the *occlusion discontinuous* case and the *self-occlusion discontinuous* case. In the continuous case, the surface patches share a common image boundary (Figure 4.5.*a*, *b* and *c*; upper boundaries). The occlusion discontinuous case results from another object occluding the object contours producing a gap between the ends of those contours (Figure 4.5.*b*, *d* and *f*). The self-occlusion discontinuous case occurs when the surface self-occludes, producing cusps and T-junctions between its contours (Figure 4.5.*c*, *e* and *f*). The difference between the occlusion discontinuity and the self-occlusion discontinuity is that, in the former, the observed T-junctions involve boundaries of different objects whereas in the latter they involve boundaries of the same object.

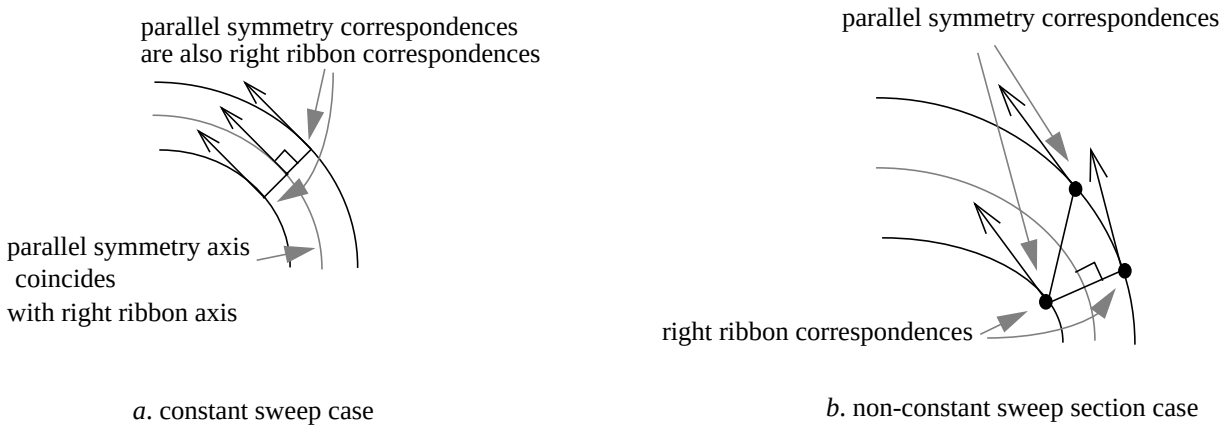

a. constant sweep case

b. non-constant sweep section case

Figure 4.4 relationship between parallel symmetry and right ribbons.

ric boundaries where the symmetric points are co-cross-section limb points (this relationship holds also for non circular PRCGCs as well [79]) and the axis of parallel symmetry the projection of the 3-D axis.

For general circular PRGCs (the non-constant cross-section case), we do not obtain exact relationships. However, QP1 implies that a right ribbon with a non-constant sweep gives a “good approximation” of co-cross-section limb points of a circular PRGC “over most viewing directions”. QP2 implies that the axis of this right ribbon also gives good approximations of the orientations of the projection of the 3-D axis tangents. The relationship between the axis of parallel symmetry and the boundaries of a general circular PRGC is not given by the above properties as P1 concerns constant sweeps only and P2 indicates that exact coincidence of the correspondences occurs *only* for constant sweeps. In the non-constant right ribbon case there is an offset, between parallel symmetry correspondences and right ribbon correspondences (Figure 4.4.*b*), which is proportional to the right ribbon sweep derivative m'. Thus, there is also an offset between parallel symmetry correspondence segments and cross-section segments of the projection of a general circular PRGC. The use of the above properties and of their relationships will be discussed in section 4.4 on hypothesizing projections of circular PRGCs.

4.2.2 Structural properties

Parts of the projection of a 3-D object are also the structural events between its boundaries. What we mean by structural events are occlusion and other junction patterns that are observed in the image. In [71] an analysis of the visibility of self-occlusion events of SHGCs was given (no reference other than a technical report is available). The analysis showed that self-occlusion occurs when the sweep derivative is higher than the tangent to the viewing angle. Projections of curved axis objects are more complex than straight axis ones due to the curvature of their axis.

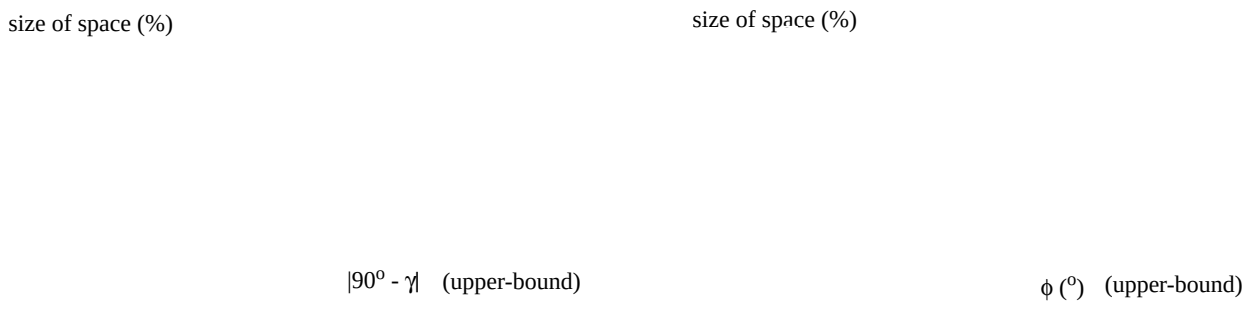

Figure 4.3 Plots of the sizes of the parameter space for different (angular) measures of properties QP1 and QP2

Invariant property IP2: In the projection of circular PRCGCs, the 2-D axis and the projection of the 3-D axis coincide exactly at corresponding points regardless of the viewing direction. Consequently, the angle, ϕ, between 2-D axis tangent and the projection of the 3-D axis tangent is zero regardless of the viewing direction.

Invariant property IP3: In the projection of a circular PRCGC, the length of the cross-section segments is constant (it is equal to the diameter of the 3-D cross-section).

Now, we give an analysis of the relationship between right ribbons and parallel symmetry. We first briefly give their definitions. See Figure 4.4.

Definition 2: A right ribbon is obtained by sweeping a straight line segment *s* along a planar axis curve *a* while scaling it by a function *m* and maintaining the axis centered at *s* and orthogonal to it. The ribbon axis is the curve *a* (Figure 4.4.*a*).

Definition 3: Two curves are parallel symmetric if there exists a continuous and monotonic pointwise correspondence function such that corresponding points have parallel tangent vectors. The axis of parallel symmetry is the locus of midpoints of the segments joining corresponding points (Figure 4.4.*a*).

Image property P1: The two sides of a right ribbon with a constant sweep function *m* are parallel symmetric with the parallel symmetric points being exactly right ribbon corresponding points (extremities of the ribbon segment *s*). Furthermore, the right ribbon axis coincides with the parallel symmetry axis (Figure 4.4.*a*).

Image property P2: The only case where the right ribbon correspondences are exactly parallel symmetry correspondences is when the right ribbon sweep derivative *m'* vanishes.

Now we relate the different properties together. Properties IP1, IP2 and IP3 imply that the projection of a circular PRCGC is *exactly* a right ribbon with a constant segment size; i.e. the right ribbon corresponding points are also co-cross-section limb points of the projection of a circular PRCGC and the ribbon axis is exactly the projection of the 3-D axis of the circular PRCGC. Property P1 together with properties IP1 through IP3 imply that the projection of a circular PRCGC produces parallel symmet-

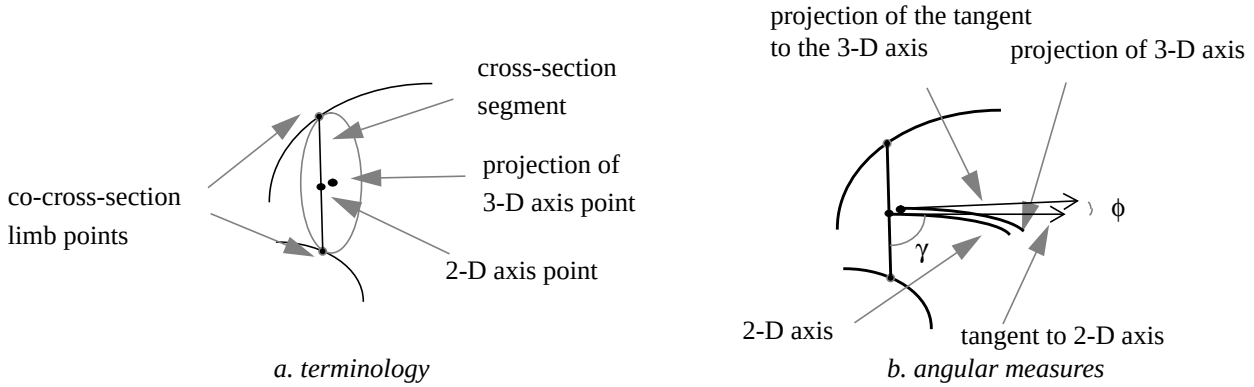

a. terminology *b. angular measures*

Figure 4.2 Properties of the projection of a circular PRGC

sures define a 4-dimensional parameter space (α, β, e, $\dot{r}$) over which the behavior of the projection of circular PRGCs can be analyzed.

Calling the projection of points on the same 3-D cross-section *co-cross-section points*, the image segment connecting co-cross-section limb points a *cross-section segmen*t and the locus of mid-points of cross-section segments the *2-D axis* (note that such mid-points are not necessarily projections of 3-D axis points; Figure 4.2.*a*), the properties are summarized here (Figure 4.2.*b*).

Quasi-invariant property QP1: In the projection of a circular PRGC, cross-section segments and the 2-D axis tangents at their mid-points are "almost orthogonal over most of the parameter space". Furthermore, the (angular) measure of orthogonality (γ) tends to degrade only close to degenerate regions of the parameter space for which limbs do not exist and which include non-general viewing directions or non-common shapes such as those close to self-intersect themselves.

Quasi-invariant property QP2: In the projection of a circular PRGC, the tangent to the 2-D axis and the projection of the tangent to the 3-D axis at the corresponding axis point are "almost parallel over most of the parameter space of observation". Furthermore, the (angular) measure of parallelism (ϕ) tends to degrade only close to the degenerate regions of the parameter space for which limbs do not exist.

The analysis to show both properties is of a statistical nature using a discretization of the parameter space. Figure 4.3 shows the plots of the sizes of the parameter space for different measures of both properties (γ is the angle of property QP1 and ϕ the angle of property QP2). For example, γ is within 5^o of 90^o over more than 84.3% of the parameter space and ϕ is 3^o or less over more than 94.48% of the parameter space. A detailed analysis of both properties is given in [82].

Properties of circular PRCGCs (constant cross-section)

Invariant property IP1: In the projection of circular PRCGCs, the image angle, γ, between cross-section segments and 2-D axis tangent is *exactly* $\pi / 2$ regardless of the viewing direction.

4.2 Projective properties of circular PRGCs

In this section we discuss the properties of the projection of a circular PRGC. The properties provide a generic object visibility model that allows the prediction of the relationships between image boundaries of the projection of a circular PRGC. They are important because objects are not given to us directly from image contours. Hypotheses about such objects have to be made, by pairing and grouping such contours, and verified based on that visibility model. To make the discussion rigorous, we start by giving a formal definition of a circular PRGC.

Definition 1: A circular PRGC is the surface obtained by (orthogonally) sweeping a circular cross-section $C(\theta)$ along a non-necessarily straight planar axis curve $A(s)$ while scaling it by a function $r(s)$.

Circular PRCGCs, the constant cross-section case, are characterized by a constant function $r(s)$. In the sub-sections below, we will review projective properties of circular PRGCs and their sub-classes. First, in section 4.2.1 we give properties that characterize the projection of the 3-D description (the geometric projective properties). Then, we discuss the structural properties in section 4.2.2.

4.2.1 Geometric projective properties

Finding the geometric properties of the projection of 3-D shapes is important to derive methods for both their detection in an image and their recovery. The properties we give characterize projections of points of the same 3-D cross-section and also relate the projection of the 3-D axis with the (2-D) axis of a particular type of ribbon commonly used in previous work (*right ribbon*). The properties also relate the projection of contour generators of circular PRGCs and another class of common symmetry (*parallel symmetry*) which we define later. We also give some (image) properties that relate right ribbons and parallel symmetry. The projective properties we discuss assume orthographic projection; i.e. that dimensions of scene objects are small compared to their distance from the camera. We start by reviewing the properties of general circular PRGCs then those of their sub-class of circular PRCGCs. The sub-class of circular PRGCs that have a straight axis (i.e. surfaces of revolution) will not be considered part of our analysis as they are also a subset of SHGCs for which segmentation and recovery methods have been derived [67,81]. For lack of space, the properties will be given without proofs.

Properties of general circular PRGCs (non-constant cross-section)

In [82], we have derived geometric *quasi-invariant* properties of general circular PRGCs. The analysis indicates that the behavior of the image contours of circular PRGCs is a function of both the viewing direction (given locally by the angles α and β in the Frenet-Serret frame along the axis) and local shape measures such as the ratio of the cross-section radius to the radius of curvature of the axis, say measure e, and the derivative, $\dot{r}$, of the scaling function. The viewing angles and the local shape mea-

Automatic segmentation and shape recovery of curved axis *3-D primitives* from a real intensity image has *not* been previously addressed in the research community to the best of our knowledge. Our approach is based on the use of two types of projective properties of such primitives. The first one, geometric projective properties, characterize the *projection of the 3-D geometric description*. They give image properties that are satisfied by projections of points on the same 3-D cross-section. It also gives strong relationships between computable 2-D and the 3-D descriptions. For example, some of them relate the axis of the projection with the projection of the axis. The second type of properties, structural properties, characterize the junctions and their relationships (structural events) that are observed in the projections of circular PRGCs. Such properties capture the complexities of the projection of 3-D shapes which, for example, include self-occlusion. Structural properties are useful in that they allow to predict and handle situations where feature discontinuities may be observed although they may be projections of continuous features in 3-D (for example, a continuous 3-D surface may project onto discontinuous 2-D surfaces due to self-occlusion).

Two types of geometric projective properties are used: *invariant* properties and *quasi-invariant* properties. Invariant properties have measures that remain unchanged with changes in viewing parameters (except perhaps for degenerate parameter values). Quasi-invariant properties have measures that vary but do so slowly and remain restricted to a relatively small range of values over a large fraction of the space of viewing parameters [7,13,82]. The invariant and quasi-invariant properties we describe give rigorous relationships between the projection of circular PRGCs and known ribbons and symmetries, a missing link in previous ribbon-based methods.

Such properties are useful in three important ways. First, they provide a *generic object visibility model* providing necessary constraints that image boundaries projecting from the same object must satisfy. Thus, they provide tools to handle contour fragmentation, surface markings and shadow boundaries. The method we describe uses them in a hypothesize-verify fashion by detecting local surface descriptions, grouping compatible ones together and verifying object hypotheses for structural consistency. Second, they provide rigorous relationships between 3-D shape and 2-D (image) properties. Thus, they can be used to generate useful constraints on 3-D shape. Finally, in making explicit use of those properties, the resulting object descriptions become independent of their particular appearance in the image.

The rest of the paper is organized as follows. In section 4.2, we discuss the properties of the projection of circular PRGCs. In section 4.3, we give an overview of the method and introduce the different steps. In section 4.4, we discuss the segmentation method in detail. Examples on real images will be given. In section 4.5, we discuss the 3-D recovery from the obtained segmented descriptions. In section 4.6, we discuss the obtained results and the robustness of the method to the parameters it uses and compare it with other methods. We conclude this paper in section 4.7.

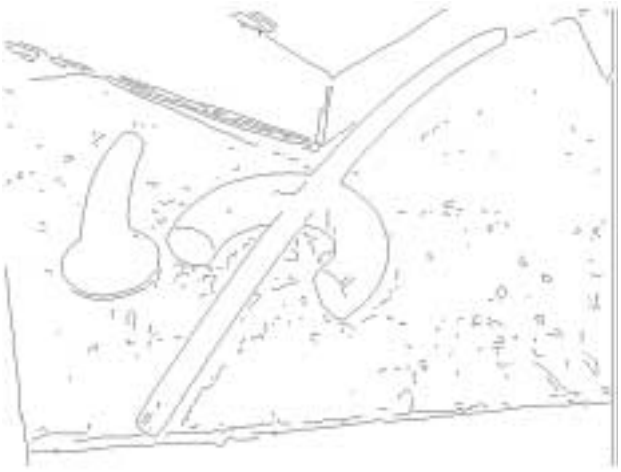

a. b.

Figure 4.1 A real intensity image (of a scene of circular PRGCs) and its extracted edges.

struments and natural objects such as snakes and animal horns. However, that method assumed that perfect and segmented boundaries were given as input. In this paper, we address the problem of segmentation and recovery of circular PRGCs from a real intensity image in the presence of occlusion, contour breaks, surface markings and shadows. Figure 4.1 shows an example of such an image.

Previous work on the recovery of GCs in the research community has shown that it is important to derive geometric projective properties of their contours. Projective invariant properties of SHGCs have been derived and used to detect partial shape information (such as axis projection) in [59], to recover 3-D shape from perfect contours [25,31,77] and to solve the figure-ground (segmentation and description) problem in real intensity images [62,81]. Curved axis primitives have been addressed in fewer efforts. In [56] mathematical derivation of limbs and cusps was discussed but without giving methods for their analysis. In [79] invariant properties of *planar right constant generalized cylinders* (PRCGCs), the constant cross-section case, have been derived and used to recover 3-D shape from perfect and segmented contours. In [82] quasi-invariant properties of circular PRGCs have been derived and used to recover 3-D shape from perfect and segmented contours.

Other related previous work includes ribbon detection methods. Ribbons have been used as the 2-D counterpart of GCs to describe image surfaces in terms of axial shape properties, including curved axes. They were used for segmentation of articulated curved objects in [50], for recognition of airplane instances in [12], for segmentation and description from imperfect contours in [62] and for detection of elongated surfaces in [34]. However, those methods did not rigorously relate ribbon descriptions to (projections of) 3-D shapes. Thus, they produce inherently 2-D descriptions which may depend on the viewpoint or occlusion patterns.

4
Segmentation and 3-D Recovery of Curved Axis Generalized Cylinders in Real Intensity Images

M. Zerroug and R. Nevatia

4.1 Introduction

Central to monocular 3-D scene analysis are the scene segmentation and shape recovery problems. Solving both problems is essential because part of describing a scene, is the description of the shape of each of its objects. However, 3-D shape recovery from a monocular image is known to be an under-constrained problem which is even more difficult for curved objects since their boundaries may be viewpoint dependent (*limb boundaries*). Furthermore, scene segmentation from a real intensity image is a particularly difficult problem because real images produce several imperfections such as noise and broken contours and surface markings, shadows and occlusion add to the difficulty.

There is great interest in solving the 3-D shape recovery from a single real intensity image since monocular images are easier (and cheaper) to acquire and, in many applications, recovering 3-D object centered descriptions is important. This includes 3-D (curved) object recognition from monocular images. In this case, the 3-D shapes offer richer and more stable descriptions than point or curve level descriptions.

We believe that it is important to develop methods of segmentation and description that handle *generic* shapes (or *classes* of shapes) rather than specific objects. This is because by addressing generic shapes, we are addressing all objects whose shapes belong to that class. Furthermore, generic shape methods have applications other than recognition including object learning and robotic grasping. The class of shapes of interest to our work are generalized cylinders (GCs) [6]. Two important sub-classes of GCs are common in our environments: straight axis and curved, planar, axis primitives. In a previous effort [81], we have developed and implemented a method for the segmentation and recovery of *straight homogeneous generalized cylinders* (SHGCs) a in real intensity image. In [82], we have derived a method for 3-D shape recovery from monocular contours of circular *planar right generalized cylinders* (circular PRGCs) which are characterized by a curved, planar, axis orthogonal to a *varying size* circular cross-section. Circular PRGCs capture man made objects such as tori and musical in-