6 From an Intensity Image to 3-D Segmented Descriptions

Mourad Zerroug and Ramakant Nevatia

We address the inference of 3-D segmented descriptions of complex objects from a single intensity image. Our approach is based on the analysis of the projective properties of a small number of generalized cylinder primitives and their relationships in the image which make up common man-made objects. Past work on this problem has either assumed perfect contours as input or used 2-dimensional shape primitives without relating them to 3-D shape. The method we present explicitly uses the 3-dimensionality of the desired descriptions and directly addresses the segmentation problem in the presence of contour breaks, markings shadows and occlusion. This work has many significant applications including recognition of complex curved objects from a single real intensity image. We demonstrate our method on real images.

6.1 Introduction

Recovering and representing shape of a complex object is one of the most fundamental tasks in computer vision. A good shape representation is useful not only for recognizing an object but also in determining how to manipulate it, how to navigate around it and to learn about new objects.

We believe that a good way to represent a complex object is by decomposing it into parts and describing the parts and the relationships between them. If the parts are complex, they can be decomposed into simpler parts and described in the same way as the larger object. Further, we believe that the parts should be described as volumetric primitives. Such a representation is very rich, stable and allows us to handle occlusion and articulation in a natural way.

Use of simpler parts to describe more complex objects has a long history in computer vision [75,83,92]. Biederman has argued that a similar scheme is used by the human visual system as well [74]. However, in spite of these theories and the obvious advantages of segmented (or part/whole) representations, their use in computer vision systems has been limited. We believe that this is due to the difficulty of actually computing segmented shape description from real data. The part decomposition hierarchy is not given in advance, we must infer it from the observable features in the data. Most of the previous work has used range data. In early work [87], Nevatia and Binford used perfect contours derived from range data. Pentland [88] used range data to segment objects into super-quadric primitives. Surface based segmentation using range data has been studied by several researchers [73,79].

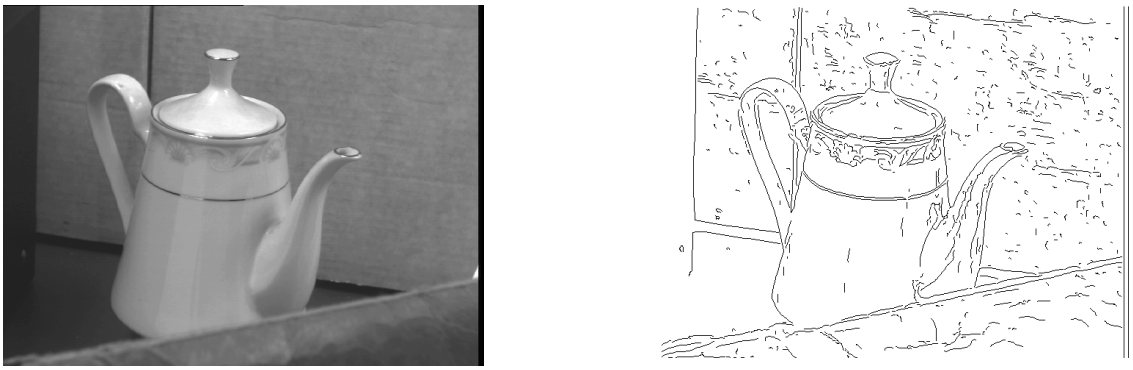

(a) intensity image (b) edge image

Figure 6.1 Sample real image of a compound object. .

In this paper, we focus on computing segmented volumetric descriptions from a single intensity image. This is a task that humans perform effortlessly. It is also important for computer vision as a single image can be acquired rather easily, without extensive control of illumination or elaborate calibration procedures. Using single intensity images does pose many problems, however. Lack of direct 3-D measurements makes it more difficult to determine discontinuities that may characterize part boundaries. Instead, we must work with intensity boundaries which may correspond to depth boundaries, but also to markings, shadows, specularities and noise. Furthermore, object boundaries are unlikely to be complete due to both poor edge localization and occlusion. These characteristics make the techniques developed for range data and perfect contours [72,77] largely unapplicable to the case of intensity images. Detection of concavities by computing curvature extrema, for example, as used in [72] is not possible; it is most likely that such extrema are missing in lines extracted from an intensity image.

Figure 6.1 shows an example. Notice that the boundaries are not all perfect, continuous or even part of the outline of the object (most are not in fact). Also, notice that the pot is partially occluded by both the spout and the flat object in front which also occludes part of the spout. The pot itself partially occludes the handle whose ends are not visible. Here, we would like to separate the teapot from the background and describe it as consisting of the arrangement of four parts: the conical pot, the lid, the spout and the handle. Deciding that there is such an object and with that composition is a non-trivial problem. Moreover, we would like to recover the 3-D shape of the object.

Some previous work has attempted to address part and object segmentation from intensity images. Specifically, the work of Rao and Nevatia [90], and Mohan and Nevatia [84] has attempted to solve similar problems to those presented in this paper. However, these efforts relied largely on heuristic properties of observed contours and did not attempt any 3-D recovery (they address a 2-dimensional problem). The method of [81], based on a neural network implementation, addresses geon-based descrip-

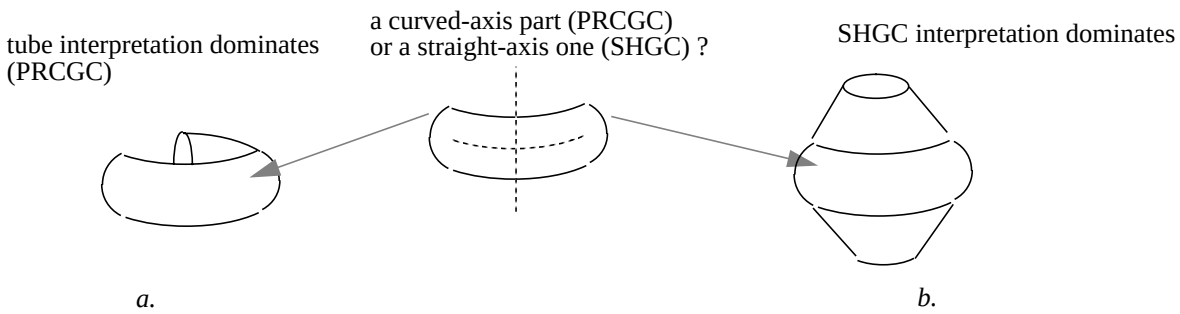

Figure 6.2 Part perception could be driven by context

tion and recognition from possibly discontinuous boundaries. The boundaries were synthetic and the axial descriptions were assumed given. In this paper, we present an approach that is more closely connected to rigorous properties of contours of 3-D objects to solve the figure/ground and part segmentation problems, *and* to recover 3-D structure of the objects.

In this work, we have chosen generalized cylinders as primitives for part description. The classes of GCs we allow here are the *straight homogeneous generalized cylinders* (SHGCs) and *planar right generalized cylinders* (PRGCs). SHGCs are obtained by scaling a planar cross-section along a straight axis curve. PRGCs are obtained by scaling a planar cross-section along a curved planar axis curve. More precisely, two sub-classes of PRGCs are addressed: *planar right constant generalized cylinders* (PRCGCs), characterized by a constant sweep, and circular PRGCs, characterized by a circular but varying size cross-section. We believe that a combination of these classes of GCs can represent well a large fraction of man-made objects.

Our approach to detecting and describing complex objects is based on the exploitation of the projective properties of the above classes of parts and of their relationships. They consist of geometric *invariant* and *quasi-invariant* and structural properties of the image boundaries of an object (the projection geometry is approximated by orthography in this work). We have used a similar approach earlier to analyze scenes of objects consisting of *single* GC primitives. Our work on SHGCs is described in [97,98,99]. Our analysis of circular PRGCs is presented in [96,97]; the method for recovering them from a single intensity image is submitted separately to this conference. Dealing with complex objects introduces many new difficulties due to the interactions between parts such as non-visibility of cross-sections, or substantial (self) occlusion, and ambiguities inherent to compound objects. For example, a part's perception depends not only on the properties of its boundaries but on the surrounding structure as well (which could be thought of as its context [71]). An example is given in Figure 6.2 . The middle part (the same set of boundaries) has different interpretations in the right and left drawings.

Our method for object segmentation and description consists of two main levels, the *part level* and the *object level*. In this approach, the figure-ground discrimination

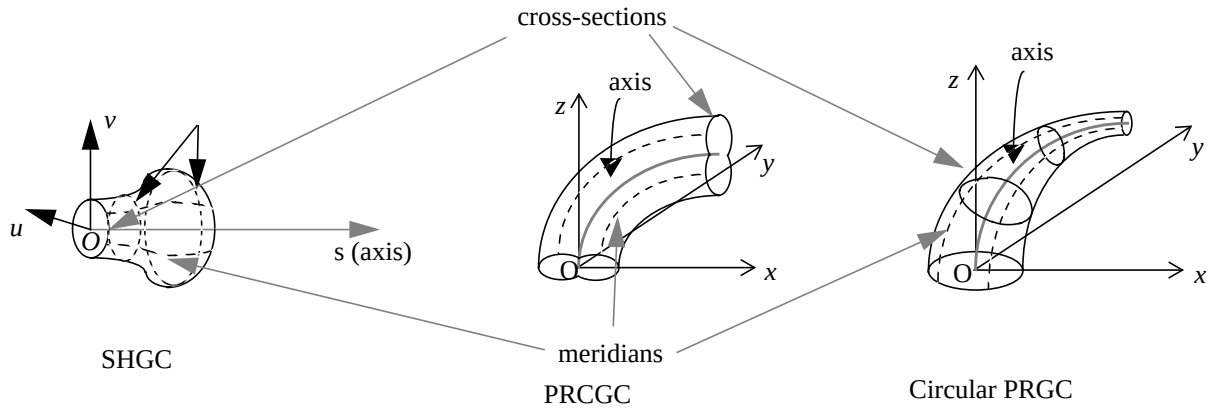

Figure 6.3 Generalized cylinders used as parts in our approach

and shape description are cooperative rather than sequential processes. The projective properties we use also help us recover the 3-D shapes of the parts we detect.

In this paper, we apply our method to a restricted (but common) class of compound objects, namely those which consist of two possible types of joints between parts: end-to-end and end-to-body. In the former, one part's end is in contact with the other part's end and in the latter, one part's end is in contact with the other part's body. Our method and some results are described in the following. First, we provide a brief overview of our part detection and description system and then describe our method for inferring the joints and the compound object.

6.2 The Part Level

Most of the methods used in this level have been described elsewhere [96,97,98,99]; due to lack of space, we will summarize them instead of giving details. The classes of parts addressed in this work (SHGCs and PRGCs) are shown in figure Figure 6.3 .

Two fundamental aspects characterize our method for detecting parts. First, it uses geometric projective (orthographic) invariant and quasi-invariant, and structural, properties of the above classes of GCs. Second it organizes the segmentation and description as a hypothesize-verify process. The projective properties provide *necessary* conditions that projections of SHGCs and PRGCs must satisfy in the image. They also give direct relationships between 3-D shapes and computable image descriptions which is useful for recovering volumetric descriptions from a monocular image. Finally, in using view invariant (and quasi-invariant) properties, the method and the descriptions it produces do not depend on the particular viewpoint the scene is viewed from.

The method for detecting parts is structured in three sub-levels: the curve level and symmetry level and the surface patch level (see figure Figure 6.4). The curve level consists of forming boundaries from image edges. The next level is concerned with

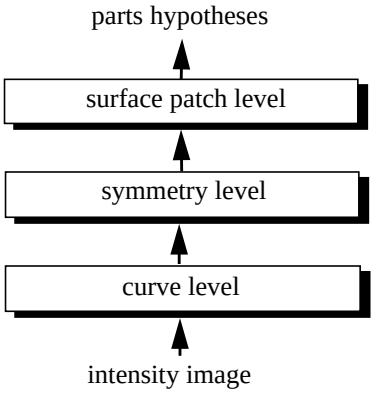

Figure 6.4 Black diagram of the part level

forming parallel symmetry [93] relationships between the boundaries (parallel symmetries are one of the invariant properties of the primitive parts). The symmetries are used to initiate the search for parts. The surface patch level is intended to form part descriptions using the boundaries and symmetries formed in the previous levels. It consists of a hypothesize-verify process of several steps: *detection* of local surface patches, *grouping* of local surface patches and *verification* of parts hypotheses. In the detection step, the projective properties are locally applied between pairs of image boundaries. Groups of boundaries which (locally) verify the properties are hypothesized to correspond to portions of parts. In the grouping step, local surface patches which are likely to project from the same scene object are merged to form parts hypotheses. The grouping criteria are based on the similarity of their projective descriptions. For example, for an SHGC, the local surface patches must have the same axis projection. The verification step consists of a filter which rules out inconsistent parts hypotheses. Consistency is defined in terms of both geometric and structural criteria. The geometric criteria consist of enforcing global consistency of the geometry of the part with respect to its geometric projective invariants and quasi-invariants. The structural criteria consist of enforcing closure and associated junctions at the end of a part. They express the fact that the image of a part may have one of several well defined closure patterns involving specific junction labeling (that include occlusion junctions) [82].

This hypothesize-verify nature of the part detection method allows us to handle markings, shadows and occlusion. Non-object boundaries (such as surface markings and shadows) are unlikely to survive the successive application of the strong projective properties. But some regular markings might still survive the verification tests.

Several enhancements, beyond our previously described work, have been made to the part level in order to handle compound objects. First, cross-sections may not be visible due to joints between parts. Second, by using a more complete set of projective properties (those of joints are described in section 6.3.1), several ambiguities occur and need to be addressed. The ambiguities are due to the fact that different 3-D

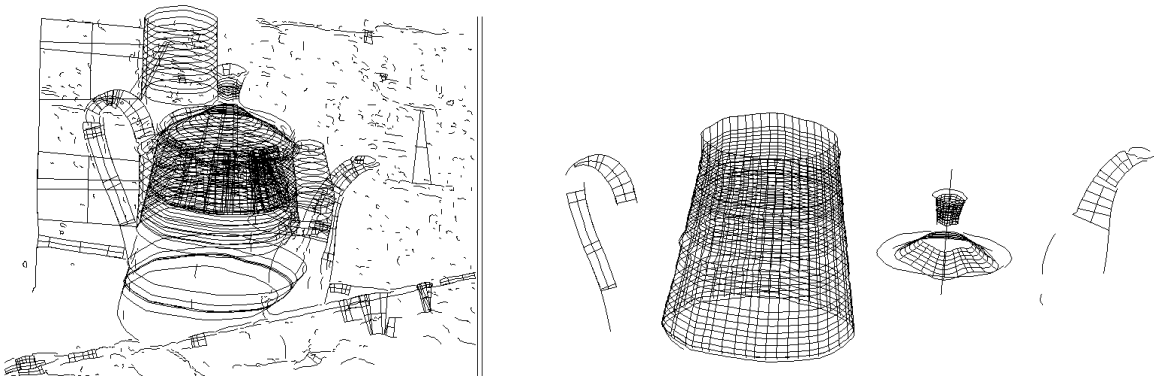

a. initial local surface patch hypotheses (146 hypotheses; only a few are shown here so as not to clutter the image)

b. verified parts hypotheses (4 hypotheses)

Figure 6.5 Results of the part level from the image of Figure 6.1 .

events could produce similar image events. For example, certain junction relationships between local surface patches could be due to self-occlusion of a single part or to a joint between different parts. For lack of space, we omit the details of the improvements to the part level.

Figure 6.5 shows results of the part level on the image of Figure 6.1 . All four verified parts consist of aggregrates of local surface patches (the pot for e.g. consists of two due to the dividing marking across its surface). Notice that the complete pot has been recovered although its boundary is occluded by both the spout and the flat object. This is discussed in section 6.6. In this example spurious hypotheses have been rejected at the verification stage.

6.3 The Object Level

Since finding complex objects consists of finding their parts and the relationships (joints) between them, it is also useful to analyze the generic image events between parts that allow us to hypothesize those relationships. In this section, we discuss the properties of the joints we address in this paper and outline the method used to form compound object descriptions.

6.3.1 Properties of Joints

There is a variety of ways parts can be joined in a compound object. In this paper, we consider two common types of joints: end-to-end and end-to-body (other types of joints could easily be incorporated). In the former the parts are joined such that their ends are in contact and in the latter such that one part's end is in contact with the other part's body (see Figure 6.6). The properties of these two types of joints consist of the closure patterns of the joined parts and the observed junction relationships between the joined parts' boundaries. Analyzing the different possibilities, as relate to

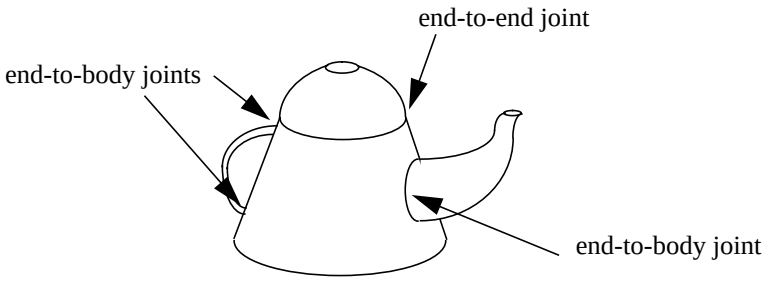

Figure 6.6 Examples of joints between parts

viewpoint for example, is useful for hypothesizing joint relationships between detected parts in the image. The properties are discussed below.

6.3.2 End-to-end joints

Our model of an end-to-end joint has two possibilities: the two cross-sections have the same size (Figure 6.7 .a through c) at their contact or have different sizes (Figure 6.7 .d and e). The observed closure patterns of the parts and the image relationships between their boundaries depend on both the parts' shape parameters (for example sweep derivatives and axis curvature) and the viewing direction (or the object's pose). In the latter case, the intersection curve (the joint curve) may or may not be visible in the image and self-occlusion may be observed. Note that the parts contact is not necessarily at their cross-sections (the joint curve may not be the cross-section curve of either part). The different arrangements for both joint closures and events between parts are shown in Figure 6.7 . The abbreviations for the junctions are as follows: *L-j* stands for *L*-junction, *T-j* for *T*-junction and *3-tgt-j* for three-tangent junction (from the catalog given in [82]).

In case *a*., there is no self-occlusion. In case *b*. there is self-occlusion and the joint curve is visible. In case *c*. there is self-occlusion and the joint curve is not visible. Cases d. and e. have self-occlusion and differ in the visibility of the joint curve. In the figure, examples with *L-j L-j* closure could be replaced by any of the image closure patterns that result from the cross-section facing away from the viewer and the examples with *3-tgt-j 3-tgt-j* closure could be replaced by any of the image closure patterns that result from the cross-section facing toward the viewer.

6.3.3 End-to-body joints

Our model of end-to-body joint consists of a part's end in contact with another part's body. The closure patterns of the joined part and the image relationships between the parts depend on whether the joint curve is visible or not. In Figure 6.8 .a, the joined part has an *L-j L-j* closure and *T*-junctions with the other part's boundaries. In Figure 6.8 .b, the joined part has *T-j T-j* closure where the *T*-junctions are with the other part's boundaries.

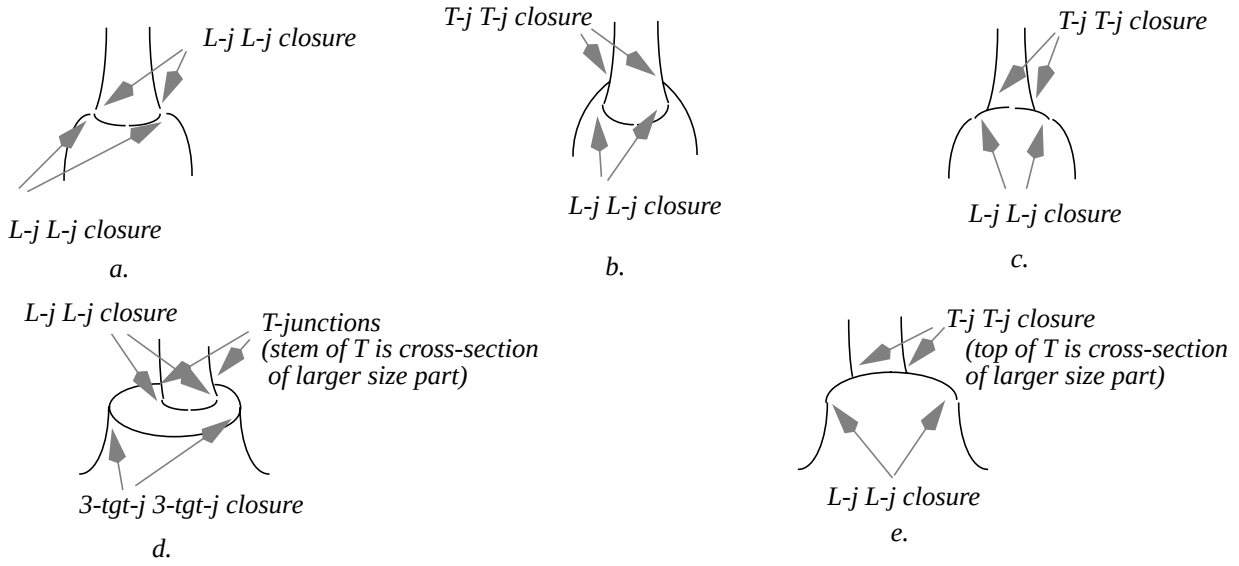

Figure 6.7 Structural relationships for end-to-end joints.
Equal size ends (a through c) and different size ends (d and e).

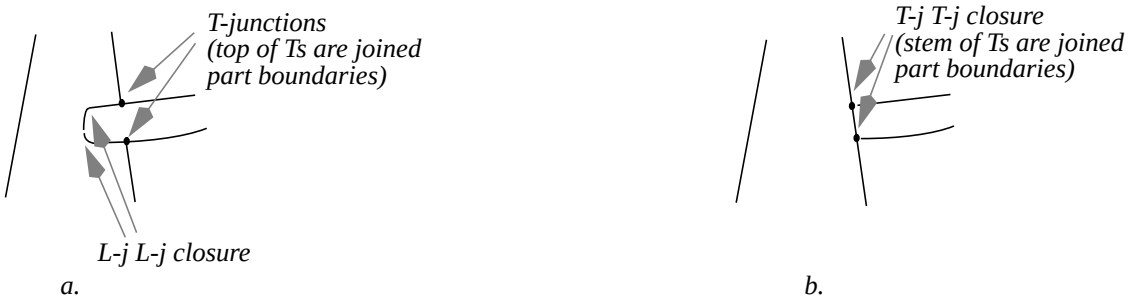

Figure 6.8 Structural relationships for end-to-body joints.
a. visible joint curve.b. non-visible joint curve

The above joint models allow for partial occlusion. Effects of occlusion by other bodies and of contour breaks are discussed in section 6.4.

6.3.4 Detection of Compound Objects

Detection of compound objects consist of hypothesizing joint relationships between detected parts. The process is not as direct as simply detecting the joint properties given previously. An inherent issue to monocular analysis of 3-D scenes is the ambiguity of the projective properties (an example was given in Figure 6.2). Thus, multiple interpretations are possible from contours alone. Further, since our goal is to produce descriptions in terms of GCs and their relationships, we must also produce descriptions that are as complete as the image allows to infer. These descriptions could be 3-dimensional if sufficient information is available in the image or otherwise 2-dimensional but corresponding to the *projections of the 3-D descriptions*. This level addresses these issues. It is organized in four steps (see Figure 6.9): detection of

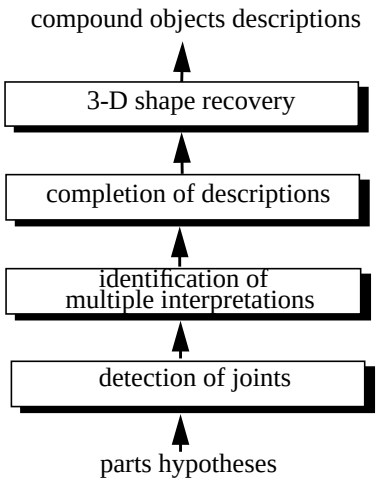

Figure 6.9 Block diagram of the object level

a. b.

Figure 6.10 Joint detection allows for partial occlusion. a. a joint is marked between parts p1 and p2 .b. no joint is marked

joints, detection of multiple interpretations, completion of description and 3-D shape recovery. We discuss these four steps below.

6.4 Detection of Joints

The objective of this step is to identify potential joint relationships between hypothesized parts. Whether there is actual (physical) contact between parts cannot be concluded from an image. The detection method consists of finding for each primitive the structural relationships of Figure 6.7 and Figure 6.8 with other primitives. Since most of those relationships involve *T*-junctions, these are first detected for all hypothesized parts (some of them are given from the analysis of the part level). The algorithm for checking any of the end-to-end or end-to-body joints between a pair of parts is fairly simple. Between a pair of parts, it uses an analysis of their contact (the same closing curve in case *a* of Figure 6.7 for e.g, or *T*-junctions between them), of the closure patterns at the joined regions and of an "extent" analysis. The latter, in case *d* of Figure 6.7 for example, consists of verifying that the closing curves of the smaller size part are all "inside" the region bounded by the cross-section boundaries of the other part's larger size cross-section. The closure constraints of the joined parts are relaxed to include occlusion at at most one side of a part's end. For example, the example of Figure 6.10 .*a* is accepted as a joint, whereas the one of Figure 6.10 .*b* is not.

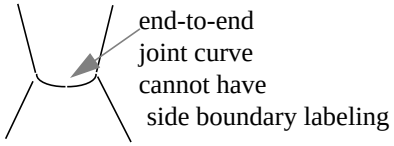

Figure 6.11 Visible joint boundaries suggest “part-end” boundary labelling

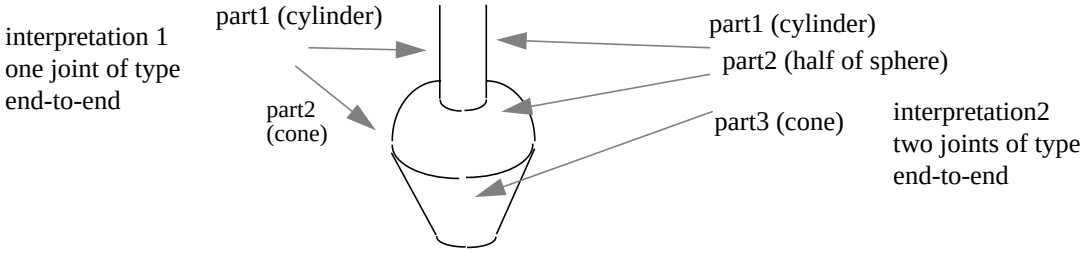

Figure 6.12 Some ambiguities persist

6.5 Identification of Multiple Interpretations

This step attempts to identify cases where more than one 3-D interpretation is possible from the given descriptions detected so far. First, the detected joints, providing global structure (or context), are used to filter out inconsistent interpretations. The joint relationships can be thought of as non-accidental relationships whose presence in the image suggests certain labelings of boundaries. The idea is that joints with visible joint boundaries (intersection curves) are unlikely to have occurred by chance in the image and they should be interpreted as parts ends. As shown in Figure 6.11 , the joint boundaries are unlikely to have side (or limb) boundary labeling (an instance of this situation was illustrated in Figure 6.2 .b). Therefore, while the single part of Figure 6.2 , taken by itself could be interpreted as either an SHGC or a PRGC, when considered in the joint of Figure 6.2 .*b*, it can only be interpreted as an SHGC part.

Remaining ambiguities are those for which certain image boundaries have different parts and joints interpretations (for e.g boundaries which could be interpreted as either cross-section or side boundaries). Figure 6.12 gives an example where two interpretations are possible: a joint of type end-to-end between two parts or two joints of type end-to-end between three parts. Therefore, conflicting parts hypotheses (two or three in this case) imply conflicting joints hypotheses (one or two joints for the same case).

The result at this stage of the method is a set of possible interpretations each of which is represented a set of graphs (one for each compound object) whose nodes are the parts and whose arcs are the joints between the parts. The arcs are labeled partially by the type of joints they represent. This graph is only a representation of the detected objects. Its purpose is not the same as the one in the method of [90] where the graph was used to segment objects made up of ribbons. Although multiple inter-

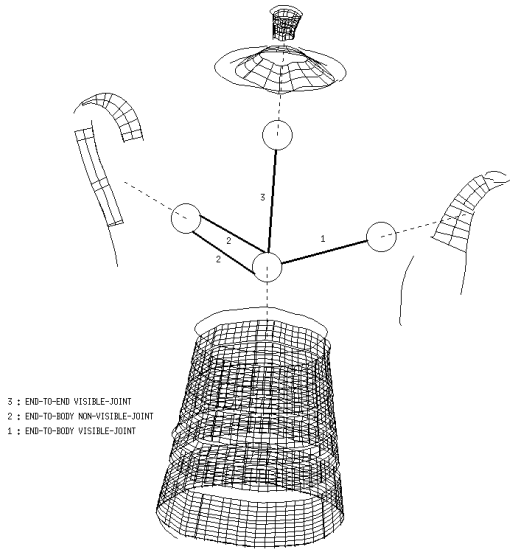

Figure 6.13 Resulting graphical representation from the hypothesized parts of Figure 6.5

pretations are a feature of our system, in the examples given in this paper, only one interpretation is found for each image. Figure 6.13 shows the graph constructed for the parts of Figure 6.5 .

6.6 Completion of Descriptions

A complete part description is one which gives its cross-section and the correspondences between its sides (projections of points on the same cross-section in 3-D), both of which give the projection of the 3-D description. Having these two elements is essential for constraining the 3-D shape of the part [93,94,95,96,97,98,99]. Some parts may already have complete projective descriptions (whole body and cross-section visible). Depending on the arrangements of parts, (self) occlusion and image contrast, some parts may miss portions of their body or their cross-sections. Two types of completions for these parts are possible: in the image and after 3-D shape is recovered.

For SHGCs with visible cross-section and partially occluded bodies, the description can be (uniquely) completed using the partial part's (projective) description [98,99]. For other parts or when the cross-section is not visible, the method attempts first to infer missing shape information, such as non-(directly)-visible cross-sections and missing side-boundaries, to the extent possible from the joint relationships between parts. For this, it is useful to classify the cut of each part as to whether it is likely to be planar or even cross-sectional. Having this classification also helps select the appropriate 3-D recovery methods [95]. To do this, we can use the geometric properties of SHGCs and PRGCs. A summary of these properties (under orthographic projection) is given below. Their use is in the reverse sense; i.e. given the observed (non-accidental) properties, we hypothesize that the part has the corresponding type of cut.

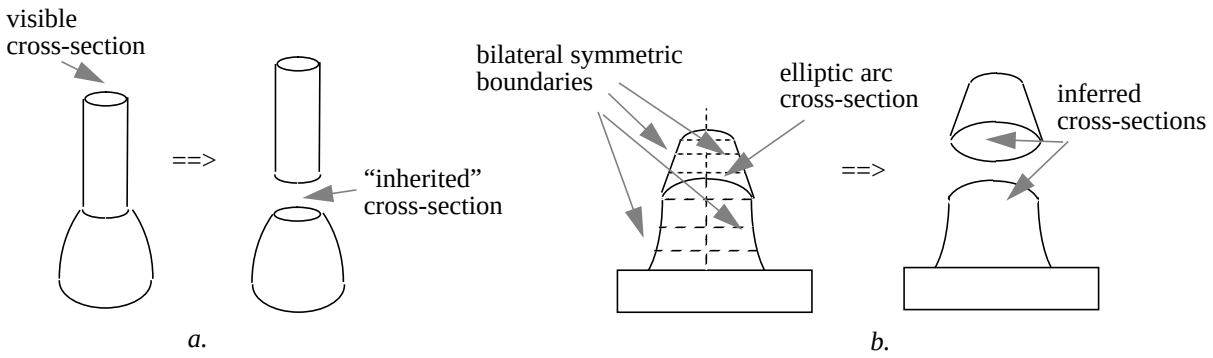

Figure 6.14 Completion of descriptions: inferring parts cross-sections

- a cross-section cut of an SHGC at both ends produces *linearly parallel symmetric*[1] boundaries in the image [93]
- a planar cut of an LSHGC at both ends produces *line-convergent*[2] symmetric boundaries in the image [95]
- a cross-sectional cut of a circular PRGC (and circular PRCGC) at any end produces an ellipse whose minor axis is *parallel* to the tangent to the projection of the axis [96]

Two types of cross-section inference are possible. One is through the use of the end-to-end joints with same-size ends (cases *a*, *b* and *c* of Figure 6.7). For example, in Figure 6.14 .a, the top part has a visible cross-section which can be "inherited" by the bottom part. This propagation can be carried through a sequence of joints. In case this is not possible, the other type of inference consists of inferring circular cross-sections for primitives which could consistently be described as such. This includes SHGCs with bilateral symmetric side boundaries [86] and circular PRGCs satisfying the third property above and having elliptic cross-sectional arcs as the partially visible cross-section. Figure 6.14 .*b* gives an example.

For parts which cannot be completed in the image and for which partial 3-D shape can be recovered, the completion is done in 3-D and consists of filling in the gaps in the 3-D axis by quadratic curves (in its recovered plane) and the gaps in the sweep function by piecewise linear sweeps.

The projective completion of the pot of Figure 6.13 (main part) has been done using only the SHGC description. The completion of the spout (right PRGC) is done after 3-D shape is recovered (see Figure 6.15). The handle (left PRGC) could not be completed since its cross-sections and joints are not visible due to occlusion.

---

1. Parallel symmetry is a generalization of parallelism of straight lines to curved ones. It is linear if the curves are scaled and translated version of each other [95].
2. Line-convergence is a form of symmetry whereby tangent lines at symmetric points of two curves intersect along a line[95]. It can be thought of as the generalization of point incidence to line incidence.

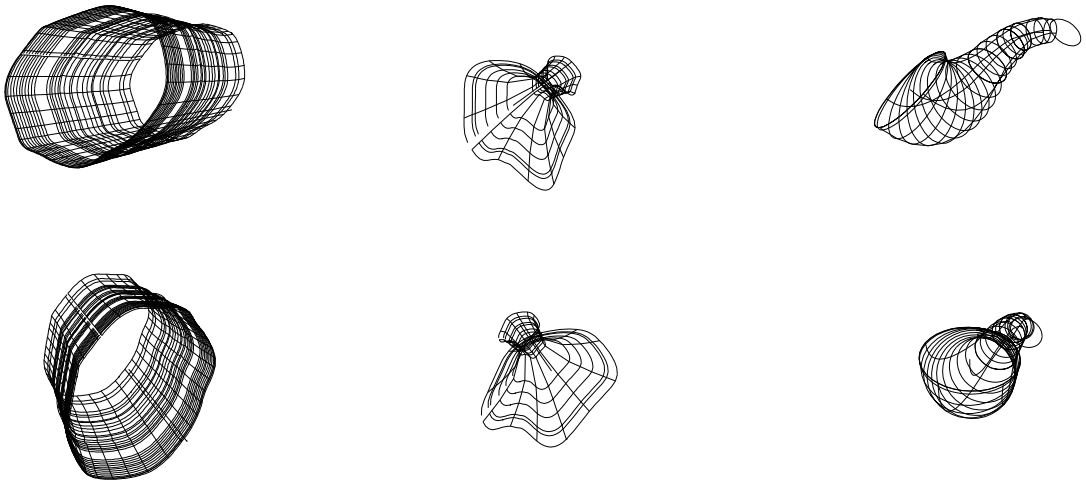

Figure 6.15 Recovered 3-D volumetric descriptions for the descriptions of Figure 6.13 .

6.7 3-D Shape Recovery

At this stage of the system, possible interpretations of image boundaries in terms of compound objects are identified. Recovering 3-D shape of a compound object consists of recovering the intrinsic 3-D description of each of its parts; i.e. its 3-D cross-section, its 3-D axis and the sweep function. 3-D descriptions from monocular images of SHGCs, PRCGCs and circular PRGCs have been addressed by the authors and others in [76,80,85,89,93,98,97], [94] and [96] respectively. Those method are based on using the properties of a primitive to generate constraints on its 3-D shape. For compound objects, a complete 3-D recovery method should normally use, besides such properties for each part, constraints on the interaction between parts. For example, well defined differential geometric relationships hold between the orientations of two surfaces and their intersection curve [78]. We have not attempted to address this problem in our work though (this is a topic in its own right).

We have instead used the methods described in [93,98,97,96] to recover 3-D shape of each part as though it were isolated. Figure 6.15 shows the recovered volumetric descriptions of the parts of the object in Figure 6.13 , using different poses in 3-D (the descriptions are shown in terms of cross-sections and meridians of the recovered 3-D GCs). Notice the 3-D completed spout. The handle could not be recovered since its cross-sections are not (even partially) visible (its description remains projective)

Figure 6.16 shows results of the method on another image. Both mugs consist of a main body (SHGC) and a handle (PRGC). Notice the markings on the surface of the cup in the front and outside the objects. The image also exhibits occlusion between independent objects. There were 177 initial local surface patch hypotheses (Figure 6.16 .c) and two objects, each made up of two parts joined by two joints, are ob-

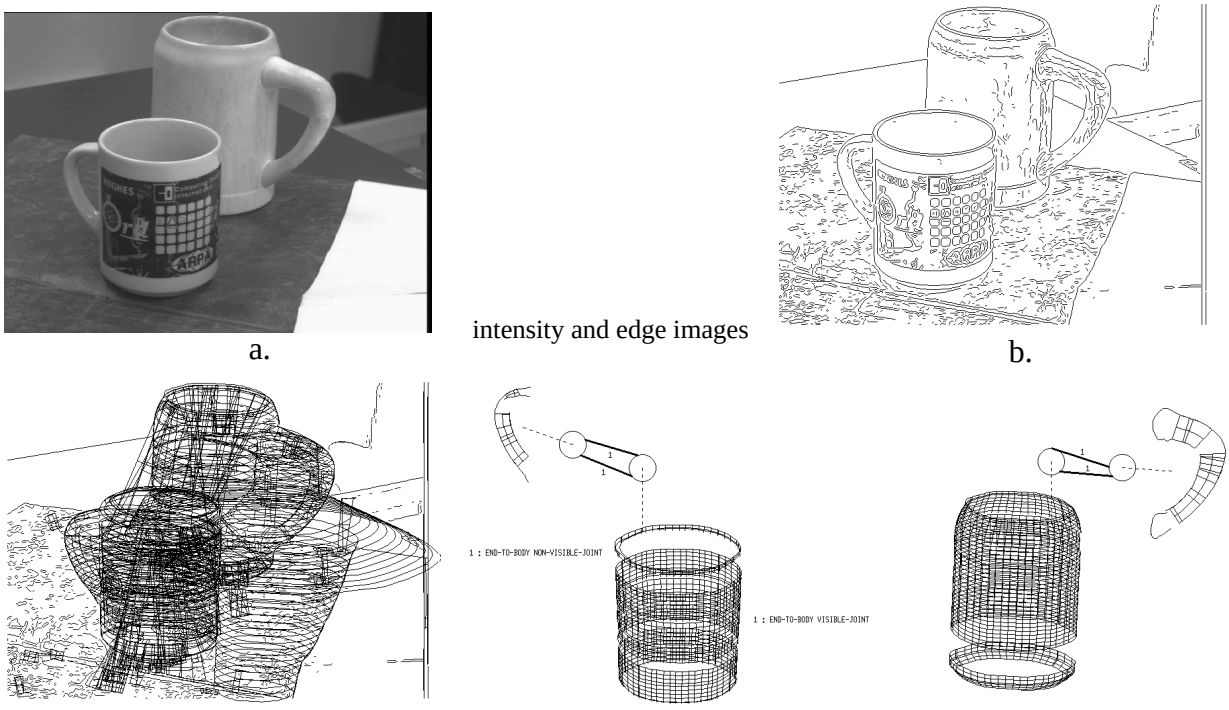

Figure 6.16 Additional results of the method

tained. The joints labeling and the obtained graphical representations are shown in Figure 6.16 .d. The recovered 3-D parts are shown in Figure 6.16 .e for different orientations. The 3-D shape of the handle of the front mug could not be recovered because its cross-sections are not visible. At this stage, the 3-D parts and their relationships are completely identified.

6.8 Discussion and Conclusion

There are a number of issues not addressed in this paper. Among them is handling parts with multiple surface patches such as occurs with polyhedral cross-section parts. This issue has been partially resolved in a previous work in the case of concave cross-section SHGCs [98,99]. To handle this type of parts requires a further step

whereby surface patches generated from different portions of the cross-section are identified and merged.

In summary, the proposed method makes use of the projective properties of a small number of primitive GCs and their relationships in order to recover segmented 3-D descriptions independently of the viewing direction and in the presence of partial occlusion, surface markings, shadows and contour breaks.

The results of this work have several applications. The descriptions obtained by our system (either the 3-D intrinsic elements of a GC or their projective descriptions) can be used to provide powerful, view-insensitive, indexing keys to large databases of object models for object recognition (such as in [87] for example). In manipulation, the 3-D descriptions can be used to plan for the grasp and pre-shape the hand. In navigation, they can be used to select appropriate paths to avoid obstacles, for example, and in learning, the symbolic descriptions can be used to analyze differences and similarities between newly recovered objects and previously recovered ones.

6.9 References for Chapter

[71] H.G. Barrow and J.M. Tenenbaum, "Interpreting Line Drawings as Three Dimensional Surfaces" Artificial Intelligence, 17:75-116, 1981.

[72] R. Bergevin and M.D. Levine, "Generic Object Recognition: Building and Matching Coarse Descriptions from Line Drawings," in *IEEE Transactions PAMI*, 15, pages 19-36, 1993.

[73] P.J. Besl and R.C. Jain, "Segmentation Through Symbolic Surface Descriptions," In *Proceedings of IEEE CVPR*, pages 77-85, 1986.

[74] I. Biederman, "Recognition by Components: A Theory of Human Image Understanding," *Psychological Review*, 94(2):115-147.

[75] T.O. Binford, "Visual Perception by Computer," *IEEE Conference on Systems and Controls*, December 1971, Miami.

[76] M. Dhome, R. Glachet and J.T. Lapreste, "Recovering the Scaling Function of a SHGC from a Single Perspective View," In *Proceedings of IEEE CVPR*, pages 36-41, 1992.

[77] S. Dickinson, "3-D shape Recovery using Distributed Aspect Matching," *IEEE Transactions PAMI*, 14(2):174-198, 1992.

[78] Do Carmo, "Differential Geometry of Curves and Surfaces," Prentice Hall, 1976.

[79] T.J. Fan, G. Medioni and R. Nevatia, "Recognizing 3-D Objects using Surface Descriptions," *IEEE Transactions PAMI*, 11(11):1140-1157, 1989.

[80] A. Gross and T. Boult, "Recovery of Generalized Cylinders from a Single Intensity View," In *Proc. Image Understanding Workshop*, pages 557-564, 1990.

[81] J.E. Hummel and I. Biederman, “Dynamic Binding in a Neural Network for Shape Recognition” *Psychological Review*, 1992.

[82] J. Malik,“Interpreting line drawings of curved objects,” *International Journal of Computer Vision*, 1(1):73-103, 1987.

[83] D. Marr, “Vision,” W.H. Freeman and Co. Publishers, 1981

[84] R. Mohan and R. Nevatia, “Perceptual organization for scene segmentation,” *IEEE Transactions PAMI*. 1992.

[85] T. Nakamura, M. Asada and Y. Shirai, “A qualitative approach to quantitative recovery of SHGC’s shape and pose from shading and contour,” In *Proceedings of IEEE CVPR*, pages 116-121, New York, 1993.

[86] V. Nalwa, "Line drawing interpretation: Bilateral symmetry," *IEEE Transactions PAMI*, 11:1117-1120, 1989.

[87] R. Nevatia and T.O. Binford, "Description and recognition of complex curved objects," *Artificial Intelligence*, 8(1):77-98, 1977.

[88] A. Pentland, “Recognition by Parts,” in *Proceedings of the ICCV*, pages 612-620, 1987.

[89] J. Ponce, D Chelberg and W.B. Mann, "Invariant properties of straight homogeneous generalized cylinders and their contours," *IEEE Transactions PAMI*, 11(9):951-966, 1989.

[90] K. Rao and R. Nevatia, "Description of complex objects from incomplete and imperfect data," In *Proceedings of the Image Understanding Workshop*, pages 399-414, Palo Alto, California, May 1989.

[91] M. Cecilia Rivara, “Algorithms for Refining Triangular Grids Suitable for Adaptive and Multigrid Techniques,” *International Journal for Numerical Methods in Engineering*, Vol. 20, pp. 745-756, 1984.

[92] L. Roberts, “Machine Perception of Three-Dimensional Solids,” MIT Press, 1965.

[93] F. Ulupinar and R. Nevatia, "Shape from contours: SHGCs," In *Proceedings of ICCV*, pages 582-582, Osaka, Japan, 1990.

[94] F. Ulupinar and R. Nevatia, "Recovering Shape from Contour for Constant Cross Section Generalized Cylinders," In *Proceedings of Computer Vision and Pattern Recognition*, pages 674-676. 1991. Maui, Hawaii.

[95] F. Ulupinar and R. Nevatia, "Perception of 3-D surfaces from 2-D contours," *IEEE Transactions PAMI*, pages 3-18, 15, 1993.

[96] M. Zerroug and R. Nevatia, "Quasi-invariant properties and 3D shape recovery of non-straight, non-constant generalized cylinders," In *Proceedings of IEEE CVPR*, pages 96-103, New York, 1993.

[97] M. Zerroug and R. Nevatia, “Using invariance and quasi-invariance for the segmentation and recovery of curved objects,” in *Proceedings of the 2nd ARPA/ESPRIT Workshop on Geometric Invariance in Computer Vision*, The Azores, 1993.

[98] M. Zerroug and R. Nevatia, “Segmentation and 3-D recovery of SHGCs from a single intensity image," to appear in *European Conference on Computer Vision*, Stockholm, 1994.

[99] M. Zerroug and R. Nevatia, “Volumetric descriptions from a single intensity image,” to appear in *International Journal of Computer Vision*.