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.

References

[1] 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.

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

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

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

[5] T.O. Binford, “Inferring Surfaces from Images,” *Artificial Intelligence*, 17:205-245, 1981.

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

[7] Do Carmo, “Differential Geometry of Curves and Surfaces,” Prentice Hall, 1976.

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

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

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

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

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

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

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

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

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

[17] 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.

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

[19] H. Sato and T.O. Binford, "Finding and recovering SHGC objects in an edge image," *Computer Vision Graphics and Image Processing*, 57(3), pages 346-356, 1993.

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

[21] 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.

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

[23] 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.

[24] 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.

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

[26] M. Zerroug and R. Nevatia, “Volumetric Descriptions from a Single Intensity Image,” to appear in the *International Journal of Computer Vision*.

[27] M. Zerroug and R. Nevatia, "Segmentation and 3-D recovery of curved axis generalized cylinders from an intensity image", Proceedings of the ICPR, Jerusalem 1994.

[28] Mourad Zerroug. “Segmentation and Inference of 3-D Descriptions from an Intensity Image”. Ph.D. dissertation, Institute for Robotics and Intelligent Systems, University of Southern California, 1994.

Currently, end-to-body joints are not used for inferring 3-D shape. Their use would require to set explicit differential geometric relationships between the joined parts' surface orientations at their intersection curve [7]. We believe this to be a research topic in its own right.

Figure 13 shows the recovered volumetric descriptions, in terms of cross-sections meridians and axes, of the parts of the object in figure 11, shown for different 3-D orientations. The pot and the lid (SHGCs) have been identified as having the same cross-section orientation, the spout (circular PRGC) has been completed in 3-D and the handle (PRCGC) could not be recovered in 3-D because its ends are not visible (its description remains projective)

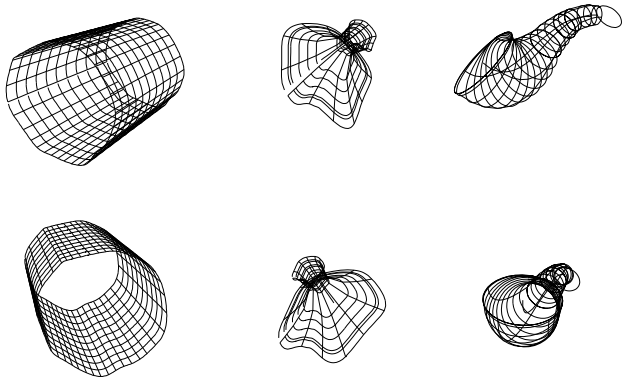

Figure 13 Recovered 3-D volumetric descriptions for the descriptions of figure 11.

Figure 14 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 front cup and on the background. The image also exhibits occlusion between independent objects. Two objects, each made up of two parts joined by two joints, are obtained. The joints labeling and the obtained graphical representations are shown in figure 14.c. The recovered 3-D parts are shown in figure 14.d 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.

a. intensity image

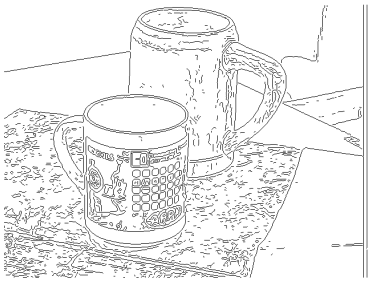

b. edge image

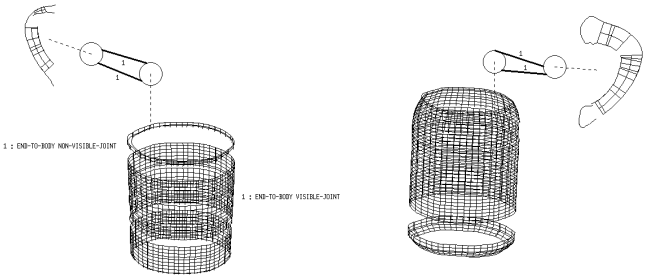

c. detected compound objects

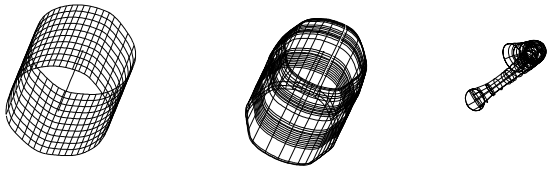

d. recovered 3-D descriptions

Figure 14 Additional results of the method

4 Conclusion

The method we described here constitutes an important progress towards realistic monocular 3-D scene analysis. It is based on a rigorous analysis of the properties of GC parts and of their generic relationships and their use to achieve both segmentation and shape description in the presence of noise, markings, shadows, contour breaks and partial occlusion.

We believe that it can be extended to handle other classes of objects and parts such as those having multiple faces (as occurs with polyhedral cross-sections for e.g.) A partial solution is given in [26].

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 [15] for example). In

3.2 Inference of 3-D shape

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. This necessitates, for each part, an image description which gives its cross-section(s) 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 [9,19,20,21,23,24,25,26]. Joint relationships between parts should also be used to add further constraints on 3-D shape.

The 3-D shape recovery of SHGCs and curved-axis primitives, as isolated parts, are discussed in [25,26] and [23,27] respectively. Here, we discuss how joints between parts can be used for 3-D shape inference. There are two such uses. The first is to set mutual constraints on the cross-section plane orientation of joined parts. The second is to infer invisible parts cross-sections.

In the former, end-to-end joints with visible joint-curve are used as indicators that the joined parts are cut by the same plane and thus their ends must have the same 3-D orientation. For this, a cut classification of the parts is useful. This classification uses evidence from the observed properties of the parts ends, to infer the type of cut (planar / cross-sectional) in an inverse manner.The classification rules are given below:

- if the part is an SHGC with linearly parallel symmetric ends then the cuts at both ends are assumed to be cross-sectional (implicitly planar) [20].
- if the part is an LSHGC with *line-convergent*[1] symmetric ends then the cuts are assumed to be planar (not both cross-sectional) [22].
- If the part is an SHGC with neither of the above symmetries, then if it has a full cross-section it is assumed to be a cross-sectional cut (for example, when the bottom end is completely occluded).
- if the part is a PRGC with at some end an extremal 2-D axis tangent orthogonal to the major axis of the elliptic cross-section then the cut is assumed to be cross-sectional (implicitly planar) [23,28].
- if the part is a PRGC with at some end elliptic cross-section not orthogonal to the extremal 2-D axis tangent then the cut is assumed to be planar but not cross-sectional.

This classification also helps select the appropriate 3-D recovery method for a part. All part ends which must have the same orientation are then identified. This is done by traversing the object graph through the arcs labeled as “end-to-end with visible joint curve” and through nodes labeled as having “parallel cuts”, such as SHGCs with linearly parallel symmetric ends. The common plane orientation of such parts ends is the average of the individual orientations obtained using the methods described in [20,21,25,26,23].

End-to-end joints with equal-size ends can also be used to infer missing parts cross-sections and consequently make their 3-D recovery possible. A missing part cross-section can be inherited from a joined part if the 3-D shape of this latter has been fully recovered. An example is given in figure 12.a. For parts whose cross-sections cannot be so inherited then circular cross-sections are assumed whenever consistent with the observed properties, including local bilateral symmetry and elliptic-arc closures [14,23,28]. Examples are given in figure 12.b (for a surface of revolution) and 12.c (for a circular PRGC).

Each time a part’s cross-section(s) is (are) so inferred, the part’s 3-D shape is recovered (as described in [25,26 ,23,27]) and the new information is propagated through the end-to-end joints with equal-size ends and visible joint-curve.

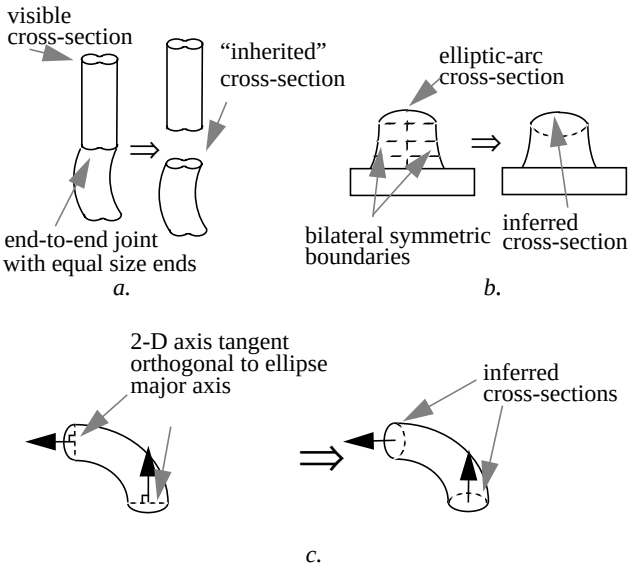

Figure 12 Completion of descriptions: inferring parts cross-sections

A 3-D shape completion step is performed for parts with gaps in their descriptions. This completion 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.

[1]Line-convergence is a form of symmetry whereby tangent lines at symmetric points of two curves intersect along a line [22]. It can be thought of as the generalization of point incidence to line incidence.

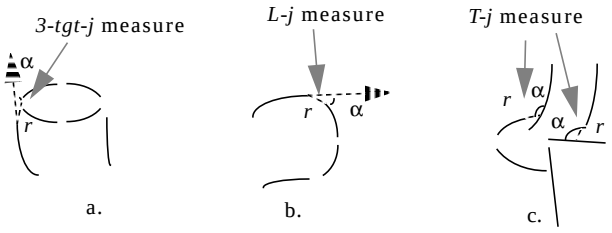

Figure 7 Junction measures.

The closure constraints of the joined parts are relaxed to include occlusion at most at one side of a part's end. For example, the case of figure 8.a is accepted as a joint, whereas the one of figure 8.b is not.

Figure 8 Joint detection allows for partial occlusion.
a. a joint is marked between parts *p*1 and *p*2;
b. no joint is marked

3.1.2 Analysis of ambiguities

Ambiguities include conflicting geometric properties (e.g. the same symmetry relationships for different GC primitives) and conflicting structural properties (e.g. same local generic events for a single part and for joined parts).

This step attempts to identify cases where more than one 3-D interpretation is possible from the given descriptions detected so far. First, since the detected joints provide a global structural context, they are used to filter out certain inconsistent interpretations. For this, joints with visible joint-curve appear to be useful. They can be thought of as non-accidental relationships whose presence in the image suggests that the boundaries of the joint-curve should be interpreted as parts ends, not sides (or limb). For example, while the single part of figure 9.a, taken by itself could be interpreted as either an SHGC or a PRCGC, when considered in the joint of figure 9.b, can only be interpreted as an SHGC part and when considered in the joint of figure 9.c as a PRCGC 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 10 gives an example where two interpretations are possible (two parts and one joint or three parts and two joints).

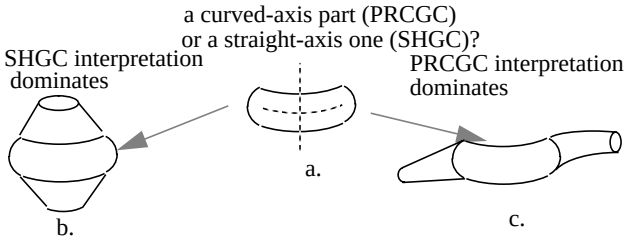

Figure 9 Part perception is linked to joint perception

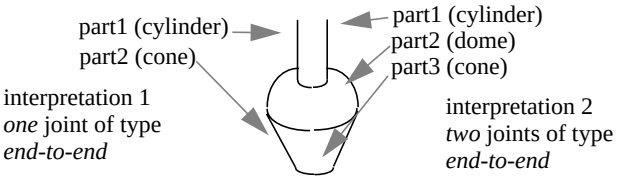

Figure 10 Some ambiguities may persist

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 [17] where the graph was used to segment objects made up of ribbons. Although multiple interpretations are a feature of our system, in the examples given in this paper, only one interpretation is found for each image. Figure 11 shows the graph constructed for the parts of figure 3.

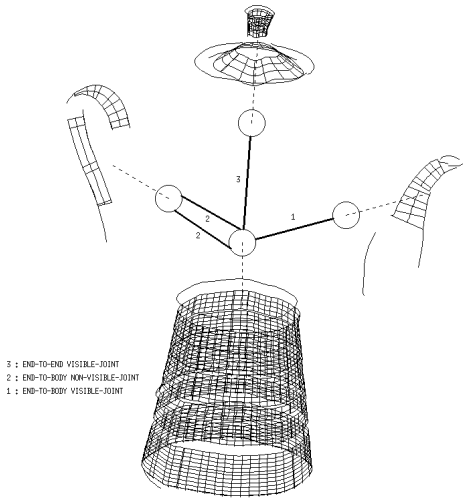

Figure 11 Resulting graphical representation from the hypothesized parts of figure 3

3.1.1 Detection of joints

The objective of this step is to identify potential joint relationships between hypothesized parts (whether there is physical contact between parts cannot be firmly concluded from an image). For this, it is useful to analyze the generic image events corresponding to the two types of joints addressed, *end-to-end* and *end-to-body* (figure 4), and use them to hypothesize those relationships.

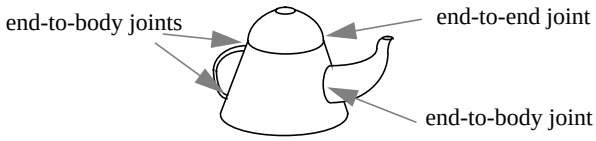

Figure 4 Examples of joints between parts

The structural 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. The observed structure depends on the viewing direction, the parts shapes (sweep derivatives and axes curvatures) at the joints and the visibility of their *joint curve* (intersection of their surfaces). In the following we only give the possible image events for each type of joint. The details of the analysis are given in [28].

End-to-end joints

Our model of an end-to-end joint has two possibilities: the two cross-sections have the same size (figure 5.a through c) at their contact or have different sizes (figure 5.d and e). The different arrangements for both joint closures and events between parts are shown in figure 5. 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 [11]; arrow and Y-junctions are currently not used). The 2-*tgt-j* corresponds to the 3-*tgt-j* case where the upper branch is not observed due to occlusion between the parts.

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.

End-to-body joints

Our model of end-to-body joint consists of a part's end in contact with another part's body. In figure 6.a, the joined part has an *L-j L-j* closure and *T*-junctions with the other part's boundaries. In figure 6.b, the joined part has *T-j T-j* closure where the *T*-junctions are with the other part's boundaries.

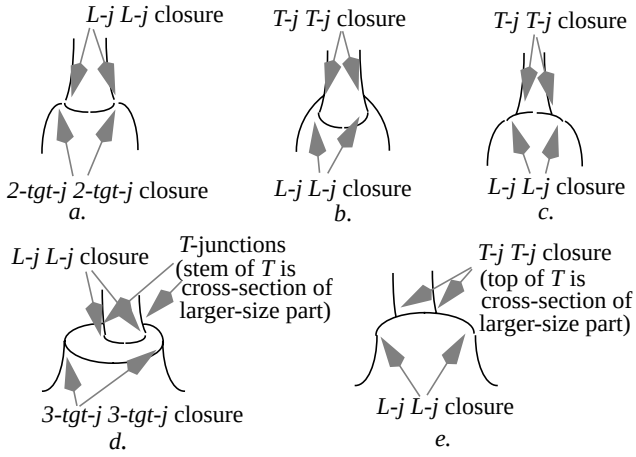

Figure 5 Structural relationships for end-to-end joints; equal-size ends (*a* through *c*) and different-size ends (*d* and *e*).

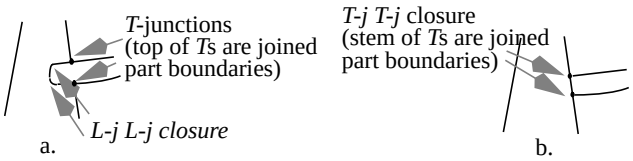

Figure 6 Structural relationships for end-to-body joints; a. visible joint-curve; b. non-visible joint-curve

In the above, 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*s the viewer (see [27]).

To detect joints, the above structural patterns are checked at the end of each part with other parts in its vicinity. The analysis, between a pair of parts, is based on the parts closure patterns, the junctions between the parts' boundaries and on an "extent" analysis. For example, the extent analysis, in case *d* of figure 5, 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 junction detection algorithms are the same as those used for the part-level verification; they make use of *junction measures* so as to account for possible boundary breaks. Figure 7 illustrates the measures for the three types of junctions used (3-*tgt-j*, *L-j* and *T-j*).

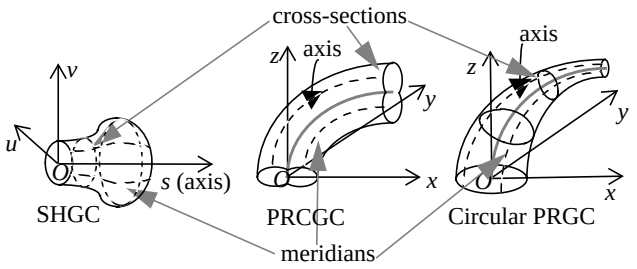

Figure 2 Generalized cylinders used as parts in our method

level 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.

The method for detecting parts consists of searching (aggregates of) boundaries that are likely to project from the same part. This is done by first *detecting* surface patches (parts fragments) that locally satisfy the projective properties of a part. Then, surface patches with "similar" projective descriptions (with respect to the expected projective properties of a part) are *grouped* into a single part hypothesis. For example, for an SHGC, the local surface patches must have the same axis projection. Each part hypothesis is then *verified* for global consistency. 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 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) [11,27]. The verification also uses an *inter-part* filtering whereby weaker part hypotheses (conflicting with "better" hypotheses) are removed.

The 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. In using the view invariant (and quasi-invariant) properties, the obtained descriptions do not depend (much) on the particular viewpoint the scene is viewed from.

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 parts and those of joints), several ambiguities occur and need to be addressed. The ambiguities are due to the fact that different 3-D events could produce similar image events (for e.g. junction patterns along a single part or between joined parts). For lack of space, we omit the details of the improvements to the part level.

Figure 3 shows results of the part level on the image of figure 1. All four verified parts consist of aggregates 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 front flat object. In this example spurious hypotheses have been rejected at the verification stage.

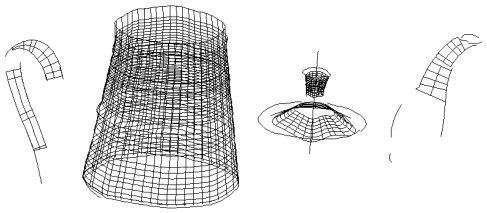

Figure 3 Results of the part level on the image of figure 1.

3 The object level

The object level uses the part hypotheses to form compound objects descriptions. The process is not as direct as simply detecting joints between parts. An inherent issue to monocular analysis of 3-D scenes is the ambiguity of the projective properties. 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 us 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 addressed these issues. It is organized in two main steps: the *detection of compound objects* and the *inference of* 3*-D shape*. The steps are discussed below.

3.1 Detection of compound objects

To detect compound objects, we have to detect joints between hypothesized parts. Because of the possible ambiguities, an explicit analysis of conflicting interpretations is also carried out. The two step are discussed below.

non-trivial problem. Moreover, we would like to recover the 3-D shape of the object.

a.

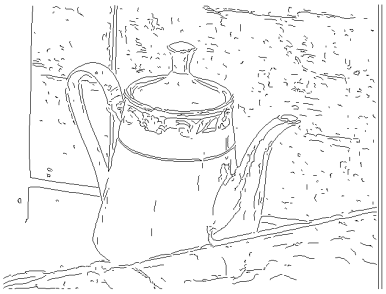

b.

Figure 1 Sample real image of a compound object. a. intensity image. b. edge image

Some previous work has attempted to address part and object segmentation from intensity images. Specifically, the work of [17] and [13] 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 [10], based on a neural network implementation, addresses geon-based description and recognition from possibly discontinuous boundaries. The boundaries were synthetic and the axial descriptions were assumed given.

1.2 Approach

In this paper, we present an approach that exploits rigorous properties of the contours of 3-D objects to solve the figure/ground (object segmentation) problems, *and* to recover 3-D structure of the objects. The class of objects addressed are those which consist of the arrangement of parts which can be described as generalized cylinders (GCs) [4]. The sub-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 planar right generalized cylinders* (circular PRGCs), characterized by a circular but varying size cross-section.

The generic joint relationships addressed include those where parts are in contact at their ends (*end-to-end* joints) and those where parts ends are in contact with other parts bodies (*end-to-body* joints). We believe that the above classes of GCs (henceforth *parts*) with such relationships 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 [24,25,26]. Our analysis of circular PRGCs is presented in [23,24]; the method for recovering curved-axis primitives from a single intensity image appears in [27]. Dealing with compound objects introduces many new issues related to the use of joint relationships to effect object-level segmentation and also shape description, including 3-D shape inference.

Our method for object segmentation and description consists of two main levels, the *part level* and the *object level*. The former is concerned with generating part hypotheses using evidence of regularity from the projective properties of parts. The latter analyzes the joint relationships between hypothesized parts to build a geometric context used to hypothesize compound objects, refine hypotheses, and infer (3-D) shape.

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 segmentation and description of compound objects.

2 The part level

Most of the methods used in this level have been described elsewhere [23,24,25,26]; due to lack of space, we will summarize them instead of giving details. The classes of parts addressed in this work (SHGCs, PRCGCs and circular PRGCs) are shown in figure 2.

Two fundamental aspects characterize our method for detecting parts. First, it uses geometric (orthographic) invariant and quasi-invariant (characterizing the intrinsic geometry of a GC), and structural, properties (characterizing lawful boundary interactions such as junctions) of the above classes of primitives. Second it organizes the segmentation and description as a multi-

Three-Dimensional Part-Based Descriptions from a Real Intensity Image*

Mourad Zerroug and Ramakant Nevatia
Institute for Robotics and Intelligent Systems
University of Southern California
Los Angeles, CA 90089-0273
zerroug@iris.usc.edu

Abstract

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.

1 Introduction

Recovering and representing the shape of a complex object is one of the most fundamental tasks in computer vision. A good shape representation is useful not only for object recognition but also for manipulation, navigation and even learning. 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. It is also highly desirable that the parts be described as volumetric primitives. Such a representation is very rich, stable and allows us to handle occlusion and articulation in a natural way.

* This research was supported in part by the Advanced Research Projects Agency of the Department of Defense and was monitored by the Air Force Office of Scientific Research under Contract No. F49620-90-C-0078 and/or Grant No F49620-93-1-0620. The United States Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notation hereon.

1.1 Previous Work

The use of simpler parts to describe more complex objects has a long history in computer vision [4,12,18]. Biederman has argued that a similar scheme is used by the human visual system as well [3]. 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 [2,8,15,16].

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 (such as based on curvature extrema) [6,1] largely unapplicable to the case of lines extracted from a real intensity image.

Figure 1 shows a sample image. Notice that the boundaries are not all perfect, continuous or even part of the outline of the object (in fact, most are not). Also, notice the partial occlusions. 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