Segmentation and Recovery of SHGCs from a Real Intensity Image[1]

Mourad Zerroug and Ramakant Nevatia

Institute for Robotics and Intelligent Systems
University of Southern California
Los Angeles, CA 90089-0273

Abstract. We address the problem of scene segmentation and shape recovery from a single real intensity image. Solving this problem is central to obtaining 3-D scene descriptions in realistic applications where perfect data cannot be obtained and only one image is available. The method we propose addresses a large class of generic shapes, namely straight homogeneous generalized cylinders (SHGCs). It consists of the derivation and use of their geometric projective properties in a multi-level grouping approach. We describe an implemented and working system that detects and recovers full SHGC descriptions in the presence of image imperfections such as broken contours, surface markings, shadows and occlusion. We demonstrate our method on complex real images.

1 Introduction

Scene segmentation and shape recovery are key problems for 3-D scene description from a monocular image. It is important to solve this problem because monocular intensity images are easier (and cheaper) to acquire and recovering 3-D descriptions is important for many applications including grasping in robotics and recognition using higher level descriptions. However, there are two major difficulties that need to be addressed. First, images provide only (2-D) *appearances* of objects and recovering 3-D shape is an under-constrained problem and a particularly difficult one for curved objects. Second, real images introduce many difficulties for a segmentation method since they usually produce broken contours, surface marking, shadows and occlusion. Figure 1.1 shows an example of a real image and its extracted edges. These imperfections make the *figure-ground* problem, a key step in scene analysis, particularly hard.

To solve this problem, we use *generic* instead of specific object models. By generic, it is meant a *class* of shapes that spans a large number of objects. In recovering generic shape descriptions, we gain higher indexing power for recognition and capabilities to learn new objects. More specifically, we address the segmentation and 3-D recovery of generalized cylinders (GCs) [3] from such monocular real images. Our approach is to handle a small number of sub-classes (primitives) of GCs: *straight ho-*

1. This research was supported 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. The United States Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notation hereon.

Lecture Notes in Computer Science, Vol. 800
Jan-Olof Eklundh (Ed.)
Computer Vision – ECCV '94

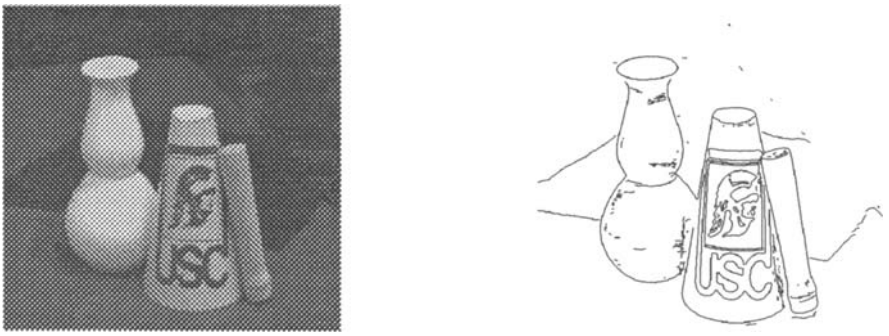

Figure 1.1 A real intensity image and its extracted edges

mogeneous generalized cylinders (SHGCs) and *planar, curved, generalized cylinders* (PRGCs). They are characterized by the scaling of a cross-section along a straight axis and a curved planar one respectively. These two primitives capture a large number of objects in our environment such as tea pots, vases and screw drivers. There is even psychological evidence that human perception of objects is strongly influenced by perception of a small number of similar primitives [2]. Since our interest is to handle compound objects made up of both types of primitives, methods of segmentation and recovery of each one of them need to be devised.

Here, we present a method for the segmentation and recovery of SHGCs in the presence of the previously mentioned real image imperfections, including occlusion. It is based on deriving and using their geometric *invariant* properties in a multi-level grouping approach. The importance of using invariant properties lies in their nature: they are viewpoint independent properties. Those that we use characterize the projection of the 3-D shape of SHGCs and provide strong constraints for their detection. This has two key consequences: the method addresses 3-D scene segmentation (produces projective descriptions) *not* image segmentation and becomes viewpoint independent which is highly desirable for 3-D scene analysis.

SHGCs and their properties have been addressed by several researchers. Their invariant properties have been addressed in [11,14,19,21,22]. Their 3-D shape recovery, including partial descriptions, has been addressed in [6,7,10,21,22]. Their recognition and pose estimation have been addressed in [8,16]. The above methods did not address the segmentation problem and full recovery of SHGCs from a real intensity image. Sato and Binford [19] did address the detection of SHGCs from a real image. Their method and ours differ in the way the properties are used and in the complexity of the scene they can handle. We will give a comparison in Sect. 5.

Most previous work on generic shape detection has either assumed perfect boundaries [1,12] and/or addressed ribbons [5,9,15] without rigorously relating the obtained descriptions to the 3-D shape of the objects they describe.

The method we present proceeds in three grouping levels: the curve level, the parallel symmetry level and SHGC patch level. The curve level uses a conservative co-curvilinearity boundary grouping and is intended to bridge short gaps and reduce the search space at subsequent levels. The parallel symmetry level is intended to hy-

pothesize SHGC cross-sections. The SHGC patch level is intended to form complete SHGC object descriptions wherever possible in the image. The constraints used in those steps are derived from the geometric projective properties of SHGCs. Resulting descriptions are then used to recover 3-D object centered descriptions in terms of the 3-D axis, 3-D cross-section and 3-D scaling function. The method and the properties we describe assume orthographic projection; i.e. objects dimensions are assumed small compared to their distance from the camera. It also assumes that scene SHGCs have (at least partially) visible cross-sections in the image and that they are not part of composite objects (addressed in a separate effort). Through the results of this work, we hope to convey a two-fold message: that invariant properties of generic shapes can be used to solve the 3-D object segmentation problem and that there is hope to solve the difficult problem of 3-D curved shape recovery *and* recognition from a single, real, 2-D image for a large class of objects.

We organize the discussion as follows. In Sect. 2 we give useful invariant properties of SHGCs. In Sect. 3 we describe our segmentation method and give examples of its application on some examples. In Sect. 4 we demonstrate the application of the resulting descriptions for 3-D shape recovery. In Sect. 5, we give a discussion about advantages and limitations of the method. We conclude in Sect. 6.

2 Properties of SHGCs

In this section, we give useful geometric properties of the projection of an SHGC. Some of them have be derived in previous work [14,20,21]. For lack of space we will omit proofs. First, we give relevant definitions.

Definition 1: An SHGC (straight homogeneous generalized cylinder) is the surface obtained by sweeping a planar cross-section curve C along a straight axis A while scaling it by a function r.

Let C be parameterized as $C(t) = (u(t), v(t))$, $r(s)$ the scaling function and α the angle between the cross-section plane and the SHGC axis, then the surface of the SHGC can be parameterized as follows (using the formulation of [20]; see Fig. 2.1):

$$S(t, s) = (u(t)\, r(s) \sin\alpha,\ v(t)\, r(s),\ s + u(t)\, r(s) \cos\alpha) \qquad (2.1)$$

When $\alpha = \pi / 2$, we obtain a *right* SHGC (RSHGC). When $r(s)$ is linear, we obtain an LSHGC. Curves of constant t are called *meridians* and curves of constant s are called *cross-sections* (also *parallels*).

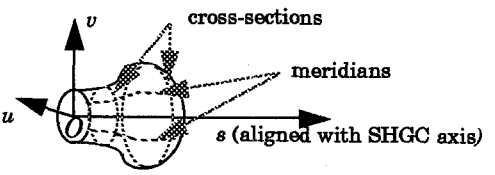

Figure 2.1 SHGC representation and terminology

Definition 2: Two planar unit speed curves $C_1(w_1)$ and $C_2(w_2)$ are said to be *parallel symmetric* [21] if there exists a continuous and monotonically increasing function f, such that $\underline{T}_1(w_1) = \underline{T}_2(w_2)$ and $w_2 = f(w_1)$. Where $\underline{T}_i(w_i)$ is the unit tangent vector of $C_i(w_i)$. That is corresponding points have parallel tangent vectors.

The correspondence is said to be linear if f is linear. In this case the two curves are similar up to scale and translation. See Fig. 2.2 for an example and terminology.

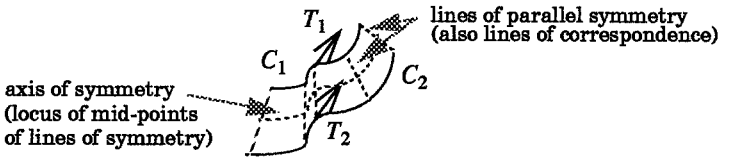

Figure 2.2 Example of parallel symmetry

Now we state the invariant properties of SHGCs. We start by those that have been derived in previous work [14,20,21] then give the new properties. See Fig. 2.3.

Property P1: Cross-section curves of an SHGC are mutually parallel symmetric with a *linear* correspondence. This property holds in 3-D and in the 2-D projection.

The proof can be found in theorem 4 and its corollary in [21].

Property P2: Contour generators (limbs) of an LSHGC are straight (they are meridians). This property holds also for the 2-D projection of limbs which are projections of those meridians. Therefore, in 2-D, the tangent line and any correspondence line at each limb point are colinear.

The proof can be found in Sect. 4 of [20].

Property P3: In 3-D, tangents to the surface in the direction of the meridians at points on the same cross-section, when not parallel, intersect at a common point on the axis of the SHGC [20]. In 2-D, tangents to the projections of limbs intersect on the projection of the axis at a common point [14,21].

The properties we add are given below. They have been mentioned, without proofs, in an overview of this work in [13]. Their proofs can be found in [25]. Equivalent properties have been independently derived by [19].

Property P4: We give this property in the form of a theorem and its corollary.
Theorem P4: *Lines of correspondence* between any pair of cross-section curves are either *parallel* to the axis or intersect on the axis at the *same point*.
Corollary P4: In 2-D, *lines of parallel symmetry* between any pair of *projected* cross-sections are either *parallel* to the *projection* of the axis or intersect on it at a *common point*, regardless of the viewing direction.
Property P5: Let $C_1(u)$ and $C_2(v)$ be two unit speed parallel symmetric curves with a linear correspondence $f(u) = au + b$. Then for all u and u' the vectors $\underline{V}_1 = C_1(u') - C_1(u)$ and $\underline{V}_2 = C_2(au' + b) - C_2(au + b)$ are *parallel* and $|\underline{V}_2| / |\underline{V}_1| = a$ (i.e. the ratio of their lengths is *constant* and equal to the *scaling* of the correspondence).

The use of these properties and the method are discussed in the next sections.

3 The Segmentation Method

The method is structured in three feature levels: the curve level, the parallel symmetry level and the SHGC patch level. For lack of space, we will not discuss the curve level and briefly discuss the parallel symmetry level in Sect. 3.1 (see [25] for a detailed account of both). We will discuss the SHGC patch level in Sect. 3.2.

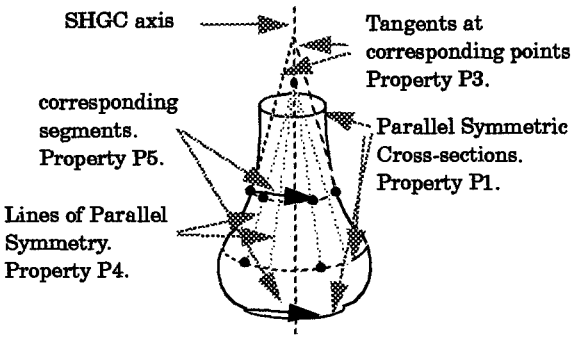

Figure 2.3 Invariant properties of SHGCs

3.1 Hypothesizing Cross-Sections (Parallel Symmetry Level)

Finding the cross-sections of an SHGC in an image allows to locate the rest of the object. We assume that such cross-sections are (at least partially) visible in the image. From Property P1, those cross-sections produce linear parallel symmetries in the image. The method to detect cross-sections consists of a hypothesize-verify process consisting of the detection of local parallel symmetry correspondences (using a quadratic B-spline representation of the boundaries [18]), grouping of compatible symmetries (thus forming global symmetries and boundaries), linearity verification of symmetries and completion of gaps (filling in missing boundaries using symmetry). The constraints used in those steps are derived from property P5.

Figure 3.1.*a* gives an example. The initial parallel symmetry correspondences are labeled ps_1 and ps_2. The grouping constraint between ps_1 and ps_2 is that the connecting segments r and s be parallel since this should be the case for an SHGC cross-section (property P5) (Fig. 3.1.b). Similarly, the (extremal) segments S and R should also be parallel. Furthermore, the length ratios $|r| / |s|$ should be equal to $|R| / |S|$ (up to some fixed threshold) (also property P5). In this case, the correspondences are linear and the gap can be completed by scaling and translating its counterpart on the symmetric curve (Fig. 3.1.c). In using the previous constraints the symmetries of Fig. 3.1.d and e are ruled out, since in case (d), the connecting segments are not parallel and in case (e) the scalings (length ratios) are not similar. Examples of results in this level are given in Fig. 3.2.

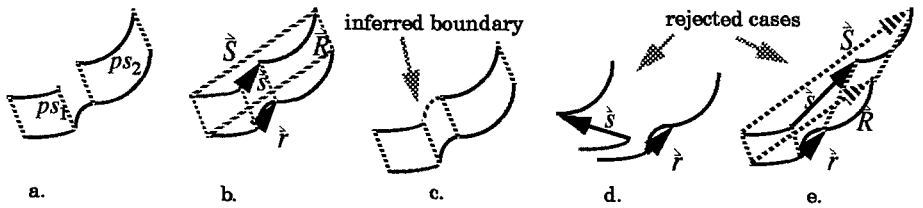

Figure 3.1 Constraints on linear parallel symmetry grouping

3.2 Detecting SHGC Descriptions (SHGC Patch Level)

The detection of SHGCs consists of *detecting* local surface patches likely to correspond to portions of SHGCs, *grouping* together those likely to belong to the

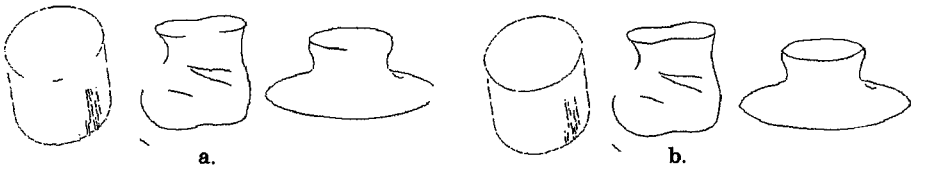

Figure 3.2 Results of the parallel symmetry level on some examples.. a. input boundaries. b. detected and completed cross-sections using linear parallel symmetry.

same object and *verifying* each object hypothesis for consistency, all using the invariant properties of SHGCs. Before describing those steps, we give the definition below.

Definition 3: A *local SHGC patch* is given by a hypothesized closed cross-section and a pair of corresponding curves satisfying the projective properties P2 or P4; i.e. they are either straight (for a local LSHGC) or have the property that lines of symmetry between any pair of projected cross-sections intersect on a straight line (projection of the axis). Figure 3.3 shows sample local SHGC patches.

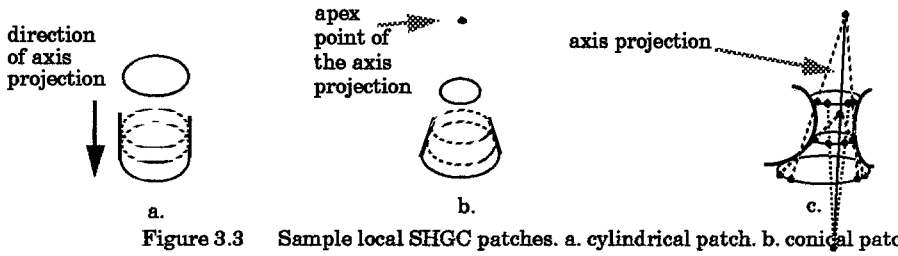

Figure 3.3 Sample local SHGC patches. a. cylindrical patch. b. conical patch. c. non-linear patch

3.2.1 Detection of Local SHGC Patches

This step consists of searching local SHGC patches for each hypothesized cross-section. For this, correspondences between each pair of curves are detected (if any). The correspondence finding method consists of ruling the area between two curves by finding the scale of the hypothesized cross-section which makes the cross-sections tangential to the limbs [21]. This yields a set of recovered cross-sections (Fig. 3.4). A pair of curves that admit such correspondences is hypothesized to form a local SHGC patch if:

- both are straight, in which case if they are parallel they form a *cylindrical patch* giving the projection of the direction of the axis; Fig. 3.3.a (corollary P4); else they form a *conical patch* whose apex is a point of the projection of the axis (Fig. 3.3.b) (also corollary P4). An LSHGC patch is thus hypothesized (property P2).
- they are not both straight; then between each pair of recovered cross-sections, the intersection point of lines of symmetry is determined (Fig. 3.3.c). A local SHGC patch is hypothesized if those points lie on a straight line (using fitting criteria). This line is a local estimate of the projection of the axis (corollary P4). We call this patch a *non-linear* SHGC patch.

This process may not result in only the "right" hypotheses (corresponding to real scene objects). There are two reasons why this is so. First, the projective invariants are necessary properties of the projections of SHGCs but *not* sufficient ones to firmly conclude their presence in the scene. Second, thresholds are used so as to ac-

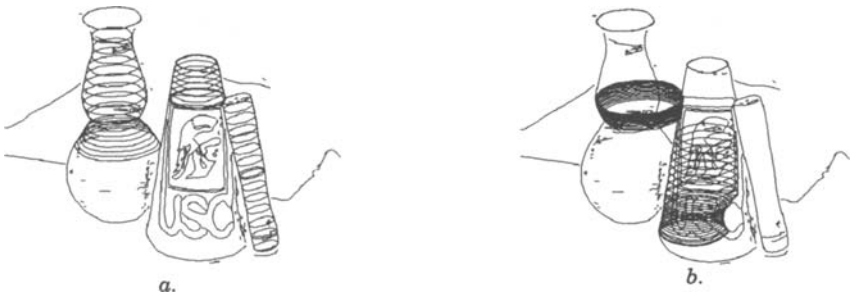

Figure 3.4 Examples of hypothesized local SHGC patches detected from the contours of Fig a. the right hypotheses. b. examples of wrong hypotheses.

count for errors in projection model, noise, quantization, etc. This inevitably results in spurious hypotheses. Figure 3.4 shows some of the local SHGC patches detected by our system from the contours of Fig. 1.1. In the figure, the "right" hypotheses are shown separately from some of the "wrong" ones, although at this stage the system cannot differentiate between them.

3.2.2 Grouping of Local SHGC Patches

Expressing the compatibility of local SHGC patches belonging to the same object is central to the grouping (segmentation) process. By a simple examination of Corollary P4, we can derive a list of viewpoint invariant geometric compatibility constraints between patches of the same SHGC. The compatibility between a pair of patches consists of the "similarity" of their local axis descriptions. The similarity relationships, depending on the type of the two patches are given below:

- *non-linear* and *non-linear*: the axes must be colinear (up to some error; Fig. 3.5.*c* and *d*)
- *non-linear* and *conical*: the cone apex must lie on the axis (up to some error; Fig. 3.5.*a*)
- *non-linear* and *cylindrical*: the direction of the cylinder must be parallel to the axis
- *conical* and *conical*: the limbs must be colinear (same apex as in Fig. 3.5.*b*) otherwise a line is generated between the apexes[2]
- *cylindrical* and *cylindrical*: the limbs must be colinear (for the same LSHGC) otherwise the directions must be parallel
- *conical* and *cylindrical*: a line from the apex in the direction of the cylinder is generated

In case, more than one grouping hypothesis at a given patch end is found, the closest one is selected. Conflicts are rare due to the strong nature of the constraints.

3.2.3 Verification of SHGC Hypotheses

In recognition methods, object hypotheses are verified by finding a transformation from model to image features. In generic shape detection, verification should

---

2. the line could be the projection the global SHGC axis; it will be later used in the verification stage.

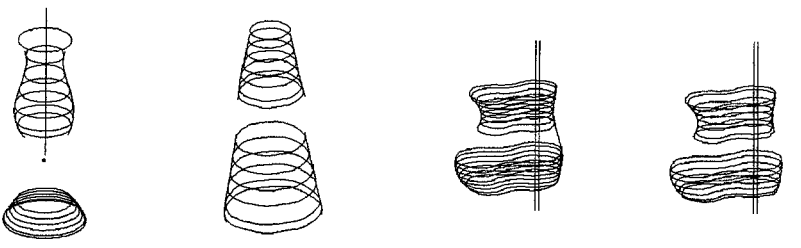

Figure 3.5 Examples of geometrically compatible local SHGC patches.

filter out the "wrong" hypotheses. The verification method we use consists of both a geometric test and a structural one. The geometric test consists of verifying that an aggregate of locally compatible patches (a global patch) is also globally compatible (the same compatibility rules as the previous ones are used). Since, the surfaces of a scene object project onto closed (or occluded) image surfaces, the structural test consists of verifying object closure using required junction properties at both ends of each hypothesized SHGC. Junction measures (based on proximity and angular variation) are used in order to account for image imperfections. However, closure may not be directly obtained due to occlusion or false negatives in the curve level grouping (see the objects in Fig. 1.1). For this, completing partial descriptions is useful. This can also be done using the projective invariants of SHGCs. Our completion method consists of two steps: the *axis-based cross-section recovery* and the *limb reconstruction* methods.

Axis-Based Cross-Section Recovery Method

This method consists of recovering cross-sections, at the gaps between patches of the same SHGC and at its ends, for continuing (unmatched) boundaries. First, for each point P_u of a given unmatched boundary, its corresponding point R_u on the reference cross-section is found (they have parallel tangents). See Fig. 3.6.a and b. The computation of the scale of the cross-section at P_u relatively to the reference one is as shown in the figure. In Fig. 3.6.a, P_c is the intersection of the line from P_u, parallel to the limb correspondence line R_u-R_c of the reference cross-section, and the other straight limb of the LSHGC. In Fig. 3.6.b, P_x is the intersection point of the line connecting P_u to R_u and the axis projection. Figure 3.7.*a* shows the cross-sections so recovered for the occluded vase of Fig. 3.4.*a* and the SHGC of Fig. 3.5.*c* .

Limb Reconstruction Method

The limb reconstruction method finds a point on each of the recovered cross-sections that is a limb point (in the projection of an SHGC, limbs and internal cross-sections are tangential to each other). The method consists of finding the *tangential envelope* of the set of recovered cross-sections (Fig. 3.6.c). This is done by approximating the space between two successive cross-sections by an LSHGC. Figure 3.7.*b* shows the limb boundaries so completed for the SHGCs of Fig. 3.7.*a*.

Selected objects are those which are both geometrically and structurally consistent (closed). Results of the whole process on the contours of Figs 1.1 and 3.2 are given in Fig. 3.8. The figure shows the global (completed) descriptions in terms of cross-sections, limbs and axes. Additional results are shown in Fig. 3.9.

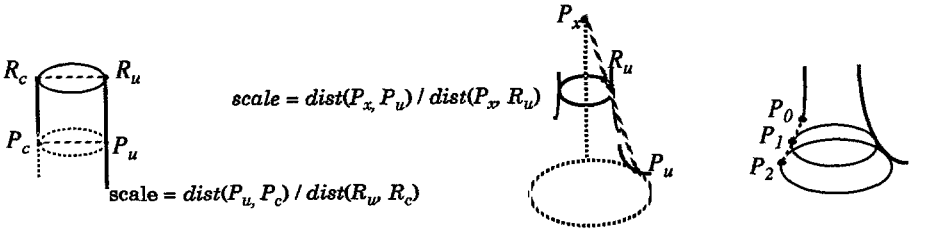

a. LSHGC b. non-linear SHGC c.

Figure 3.6 Axis based cross-section recovery (a, b) and limb reconstruction (c).

a. b.

Figure 3.7 Cross-section recovery and limb reconstruction for previous SHGCs.

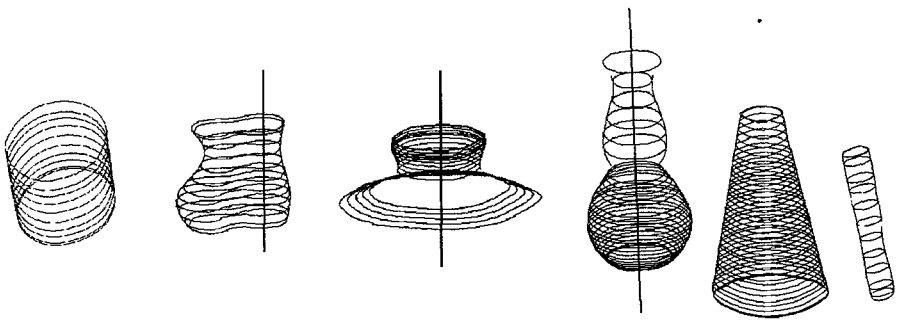

Figure 3.8 Results obtained on previous contours. The last example is from the image of Fig

4 3-D Shape Recovery

To recover object-centered descriptions, consisting of the 3-D cross-section curve, the 3-D axis and the 3-D scaling function, we build on the viewer-centered method of Ulupinar and Nevatia [21]. For lack of space, we omit the mathematical formulation, which can be found in [25]. Results of the recovery method on the descriptions of Figs. 3.8 and 3.9 are shown in Fig. 4.1. The figure shows the recovered primitives for different 3-D poses.

5 Discussion

The method is not limited to special types of cross-sections nor does it assume bilaterally symmetric boundaries. The second (from the left) SHGC of Fig. 3.8 is an example of non-circular cross-section SHGC, as is the occluded object in Fig. 3.9. The method also handles concave cross-sections which produce several self-occluding limb boundaries in the image. A *merging* step is performed on the obtained patches which have the same geometric description. Figure 5.1 illustrates this situation us-

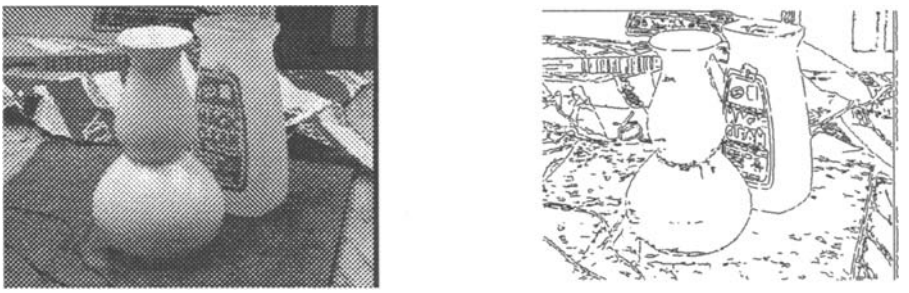

a. intensity and edge images

b. detected local SHGC patches
(only the "right" hypotheses are shown)

c.completed descriptions
of the segmented objects

Figure 3.9 Additional example of results obtained by our method.

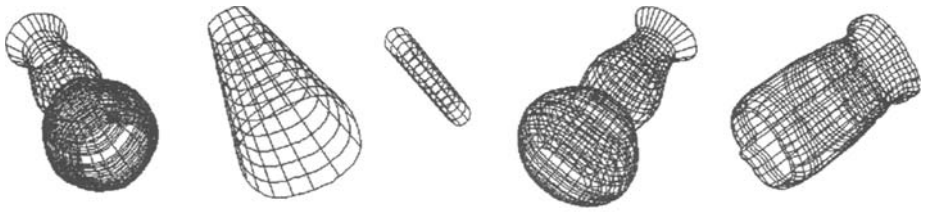

a. b.

Figure 4.1 Recovered 3-D descriptions of previous SHGC scenes shown for differen poses. a. from Fig. 3.8. b. from Fig. 3.9.

ing a synthetic object. The method also handles cross-sections with tangent discontinuities (see [25] for a more detailed discussion on such cases).

The method has been applied to several images, about 15, including variations of the ones shown in this paper. Robustness of the system to changes in the parameters it uses has been tested by changing their values by 50% of their default ones. Those changes have only affected the number of hypotheses (for e.g 95 local patches instead of the original 94 in the case of Fig. 1.1). Most importantly, the same final results have been obtained. This is due to the strong nature of the constraints.

Compared to the method of [19], our method is similar in the principle of using SHGC properties for their detection. It differs, however, in two ways. First, application of the properties in their method is somewhat restricted to surfaces of revolution and LSHGCs. Their application of the properties to surface detection, for example, may not give accurate results for general SHGCs as limb projections are generally *not* meridian projections. The main difference between their method and

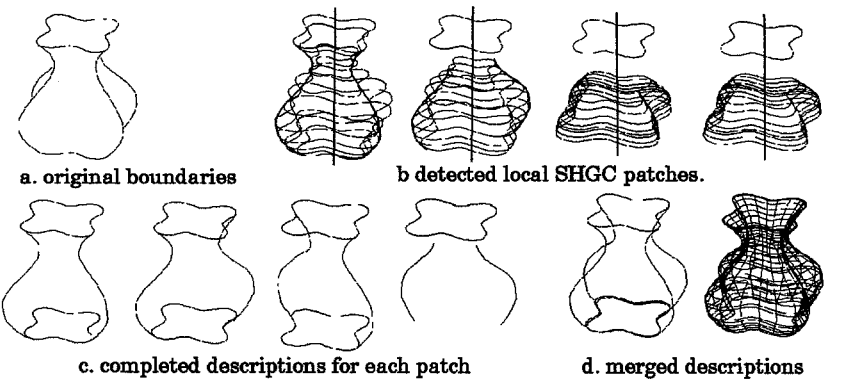

Figure 5.1 Illustrative result on a concave cross-section SHGC with multiple (merged) hypotheses

ours, however, lies in handling occlusion and large gaps. The authors note that their system does not handle occlusion as boundary connectivity criteria are used for surface detection. The detect-group-verify nature of our method allows to handle occlusion and contour breaks and even complete resulting partial descriptions.

6 Conclusion

This work demonstrates that geometric projective properties of generic shapes are a promising tool for recovering 3-D descriptions from a single real image that includes noise, broken contours, shadows, markings and occlusion. The type of properties we have presented relate computable 2-D descriptions to (the projection of) 3-D descriptions of generalized cylinders. An important characteristic of the properties is that while they are of a global nature, capturing stable and rich properties of higher level features such as surfaces, they can be locally applied to visible parts of those features, thus handling image discontinuities such as occlusion. A major consequence of using invariant properties of shape is that scene segmentation and shape recovery become viewpoint independent processes, a crucial feature of any method of generic shape description and recognition.

The results of the method have several applications including recognition. Since the shape information is reduced to the cross-section curve, the axis and the scaling function, each of these entities can be used to derive indexing keys, possibly based on qualitative measures, to a data base of object models so represented.

The method we have described has its limitations, however. Since it is based on hypothesizing presence of SHGCs using evidence from the invariant properties, it fails when such properties are not observed in the image. For example, cross-sections may not be visible in the case of compound objects due to joints between parts. Such difficulties are currently being studied.

In another effort [23,24], we have addressed a class of curved axis primitives (circular PRGCs) and derived geometric *quasi-invariant* properties and showed their usage to recover 3-D shape from monocular contours. This work and the one in [23,24] constitute the basis of a more general method, currently in development, for handling compound objects made up of both straight and curved axis primitives.

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] I. Biederman, "Recognition by components: A theory of human image understanding", *Psychological Review*, 94(2):115-147.

[3] T.O. Binford, "Visual perception by computer," *IEEE Conference on Systems and Controls*, December 1971, Miami.

[4] T.O. Binford, "Inferring surfaces from images," *Artificial Intelligence*, 17:205-245, 1981.

[5] R.A. Brooks, "Model-based three dimensional interpretation of two dimensional images," *IEEE Transactions PAMI*, 5(2):140-150, 1983.

[6] 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.

[7] A. Gross and T. Boult, "Recovery of generalized cylinders from a single intensity view," In *Proceedings of the Image Understanding Workshop*, pages 557-564, Pennsylvania, 1990.

[8] J. Liu, J. Mundy, D. Forsyth, A. Zisserman and C. Rothwell, "Efficient recognition of rotationally symmetric surfaces and straight homogeneous generalized cylinders," In *Proceedings of IEEE CVPR*, pages 123-128, 1993.

[9] R. Mohan and R. Nevatia, "Perceptual organization for scene segmentation", *IEEE Transactions PAMI*. 1992.

[10] 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.

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

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

[13] R. Nevatia, K. Price and G. Medioni, "USC image understanding research: 1990-1991". In *Proceedings of the Image Understanding Workshop*, San Diego, California, 1991.

[14] 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.

[15] 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.

[16] M. Richetin, M. Dhome, J.T. Lapestre and G. Rives, "Inverse Perspective Transform Using Zero-Curvature Contours Points: Applications to the Localization of Some Generalized Cylinders from a Single View," *IEEE Transactions PAMI*, 13(2):185-192, 1991.

[17] C.A. Rothwell, D.A. Forsyth, A. Zisserman and J.L Mundy, "Extracting projective structure from single perspective views of 3D point sets", in the proceedings of the *ICCV*, pages 573-582, Berlin, Germany. 1993.

[18] P. Saint-Marc and G. Medioni, "B-spline contour representation and symmetry detection," In *First ECCV*, pages 604-606, Antibes, France, April 1990.

[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] S.A. Shafer and T. Kanade, "The theory of straight homogeneous generalized cylinders," Technical Report CS-083-105, Carnegie Mellon University, 1983.

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

[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 International Workshop on geometric invariance in computer vision*, The Azores, 1993.

[25] M. Zerroug and R. Nevatia, "Volumetric descriptions from a single intensity image," to appear in IJCV.