

Figure 6.8 Results on tested objects obtained in level 2

7 SHGC Patch Level

At this level, the objective is to produce SHGC descriptions using curves and symmetries obtained in the previous two levels. Due to large gaps, usually caused by occlusion or simply weak contrasts, a single surface may not be detected simply by searching for closed contours, or by expecting connectivity between surface extremities as in [Sato & Binford 1992a]. Furthermore, junctions and corners may not be reliable as cues for surface segmentation in a real image, as those

- a) boundaries are inferred up to the extremities of the continuing curves (dashed curves in Figure 6.3.b). The same procedure as the one discussed previously is used here.
- b) the two remaining gaps are filled in each by a quadratic B-spline (dotted curves in Figure 6.3.b).

The two filled in boundaries are parallel symmetric (since they are obtained by quadratic B-spline segments having mutually parallel tangents at their extremities), thus producing a consistent boundary completion with the symmetry requirements. Figure 6.8.c is an example of such completion.

Finally, only symmetries involving closed curves are selected. Closure is defined by the existence of a cycle of curves connecting both extremities of a given curve. Gaps between adjacent curves in the cycle are completed by B-splines. This method has produced satisfactory results for all tested examples. Closed curves involved in parallel symmetries are likely to correspond to cross-sections of SHGCs (Property P1).

In this step, among competing hypotheses, very few survive the verifications as the constraints are rather strong. In case there remains competing hypotheses, they involve two types of situations: many *top* (closed) cross-sections for one *bottom* (other end) and many *bottoms* for one *top* (Figure 6.7 shows an example of many tops to one bottom). Each of those hypotheses would yield a different interpretation. The selection of these latter requires more global criteria (such as surface closure). Global selection is thus made at the end of the following level which we discuss next.




Figure 6.7 *Competing hypotheses may yield different interpretations*

Figure 6.8 shows the results obtained in this level on some objects. Figure 6.8.d shows the completed cross-section of the cone of Figure 5.2.b




Figure 6.6 Verification of global correspondences

where $\|p_m - p_0\| / \|q_m - q_0\|$ is the *global scale* of the compound correspondence and $\|p_j - p_i\| / \|q_j - q_i\|$ is the *local scale* of any of its components.

6.5 Boundary Completion

Selected global correspondences are used in order to fill in the gaps. Since symmetries are similarity relationships, missing boundaries of a curve can be inferred from corresponding boundaries of a symmetric curve. Boundary completion is different for the two types of connections discussed in 6.2. We discuss them separately.

1) Continuous connection

In this case, the common curve of the connected symmetry elements is used as a *model* for the missing boundary of the connection. This is done as follows.

- a) the part of the model curve that corresponds to the gap is detected. For this, the global scale is used.
- b) the missing boundary is obtained by scaling and translating the previous part so that it fills the gap.

This is shown by the dashed curve in Figure 6.3.a. This operation is done efficiently by the use of B-splines. The cross-section gaps of Figure 6.8.a and b have been so completed.

2) Discontinuous connection

In this case, there are gaps on both sides of the connection. The completion is done in two steps.

ing hypotheses. These latter are handled as follows. First, *conflict sets* that include mutually competing hypotheses are constructed. Then, within each set a simple filtering is performed that discards redundant grouping hypotheses. These latter are defined as connections between symmetry elements, say ps_i and ps_j , for which there exists a sequence of ‘atomic’ connections between the sequence ps_i ps_{i+1} ... ps_j . In other words, a sequence of short connections is preferred over one large connection since shorter connections involve more of the object boundaries.

Among the remaining hypotheses within each set, it is difficult to decide at this level which (if any) is the right one. Consequently, all alternative combinations involving non competing hypotheses (conflict free sets) are investigated and verified for consistency. This is discussed in the next section.





Figure 6.5 Competing hypotheses. Grouping of ps_1 and ps_2 competes with that of ps_3 and ps_4 .

6.4 Verification of Global Correspondences

In this step, hypothesized connections are checked for *geometric consistency*. The objective is to retain only those groupings that produce global parallel symmetries with *linear* correspondences. The verification consists of checking the similarity between the scale given by the global correspondence and the scales of each of its component parallel symmetry elements and connections (property P5). This is necessary because the local compatibility constraint only ensures scale similarity of two neighboring local correspondences (due to similarity measures, the relationship may not be transitive). Global verification is done by selecting correspondences that satisfy the following requirement (see Figure 6.6):

$$|1 - ((\|p_j - p_i\| / \|q_j - q_i\|) / (\|p_m - p_0\| / \|q_m - q_0\|))| < \varepsilon \quad \forall ij \quad (6.1)$$




$$\text{global-scale} = \frac{|\hat{S}|}{|\hat{R}|}$$

$$\text{local-scale} = \frac{|\hat{s}|}{|\hat{r}|}$$

$$E = \frac{|\hat{s}|}{|\hat{S}|}$$


a. Continuous Connection




$$E = \frac{|\hat{s}|}{|\hat{S}|} (\alpha^2 + \beta^2)$$

b. Discontinuous Connection

Figure 6.3 Grouping of parallel symmetry elements



a.



b.

Figure 6.4 Non grouped symmetries. a. Non parallel connections. b. Non similar local and global scales.

6.3 Selection of Hypotheses

The previous step may produce a large number of connection hypotheses, some of them conflicting. Conflicts arise when there is more than one connection hypothesis involving the same curve at the same end. Figure 6.5 gives an example. A final interpretation may not include conflict-

For this, we propose a perceptual grouping method based on a local compatibility constraint derived from property P5. As shown in Figure 6.3, two symmetry elements ps_1 and ps_2 are considered for grouping if the vectors \vec{s} and \vec{r} defined by the end points of the symmetries are parallel and the ratio of their lengths $|\vec{s}|/|\vec{r}|$ (local scale) is similar to the scale $|\vec{S}|/|\vec{R}|$ suggested by the grouped symmetries (global scale). In practice, we also use a connection measure for each grouping hypothesis. For this, we distinguish two cases, continuous connection and discontinuous connection.

1) Continuous Connection

In this case the two symmetry elements have a common curve (Figure 6.3.a). The local compatibility constraint, in addition to the requirements mentioned above, involves a simple connection measure based on the gap between the non connected curves. As shown in the figure, the quantity $E = |\vec{s}|/|\vec{S}|$ measures the relative length of the gap with respect to the symmetric curves. A grouping hypothesis is generated if E is less than a fixed threshold. This measure has been introduced so as to penalize distant symmetry elements, a case that often occurs between unrelated symmetries (involving markings, for example).

2) Discontinuous Connection

In this case, the symmetry elements do not share a common curve. This case can happen, for example, due to occlusion that hides parts of both end cross-sections of an SHGC. Another connection measure is used in this case, $E = (|\vec{s}|/|\vec{S}|) (\alpha^2 + \beta^2)$, and is as shown in Figure 6.3.b. It controls both the relative gap and the continuity of the symmetric curves. Gaps that involve a change in the sign of curvature are not considered.

This local compatibility constraint prevents grouping of symmetry elements such as the ones of Figure 6.4.a and Figure 6.4.b. In the former, the connecting segments are not parallel and in the latter, the local scale is not similar to the global one.

Notice that grouping of symmetry elements implies grouping of curves involved in the symmetries. Parallel symmetries provide a more global grouping criterion than co-curvilinearity. Therefore, this step implicitly handles (cross-section) curve groupings that have not been generated at the curve level due to large gaps or non smooth connections.




Figure 6.1 Block diagram of the parallel symmetry level




Figure 6.2 Local parallel symmetry correspondences (axes are in thick lines)

6.2 Grouping of Parallel Symmetries

The purpose of this step is to generate grouping hypotheses of related symmetry elements.

However, not all related curves are necessarily grouped, as the decisions are affected by the choice of the thresholds and as large gaps can exist due, say, to occlusion. Further, some grouping hypotheses may be erroneous. The next two levels can correct such errors. This is discussed in the next two sections.

6 Parallel Symmetry Level Grouping

As discussed in section 3, parallel symmetries are a projective invariant property of the cross-sections of an SHGC. They provide strong relationships between such curves in the image. Thus, their detection is an important part of the process of hypothesizing presence of SHGCs in a scene.

Detection of parallel symmetries has been investigated by few researchers. Saint-Marc and Medioni [1990] have proposed a method based on B-spline representation of image curves. That method has the advantage of being efficient (small number of B-spline segments). However, it may not produce only desired symmetries particularly in the presence of large gaps that were not bridged in the first level.

Our method consists of using the local correspondences produced by that method and a hypothesize-verify process to *group* relevant symmetry elements into more global correspondences. Several steps make up this level as shown in the diagram of Figure 6.1. We discuss each step separately.

6.1 Detection of Local Correspondences

The detection method of [Saint-Marc 1990] consists, first, of fitting quadratic B-splines to curves then finding correspondences analytically. All possible pairings between curves are considered in order to produce all such correspondences (we omit the details of this method as they are described in the previous reference). The correspondences obtained are generally noisy, sparse due to breaks and may involve both desired and undesired symmetries (involving markings, for example). Further, the correspondences may not be linear, which is a requirement for cross-sections of SHGCs (property P1). Figure 6.2 gives some examples of correspondences given by that method. The second example has of the order of a thousand such symmetries. Grouping and selection of relevant correspondences are the objective of the next steps.

level and consists of defining a local compatibility constraint based on an energy measure between curve-ends. This energy measure is $E = [l / (l_1 + l + l_2)] [\alpha^2 + \beta^2]$ where the variables are as shown in Figure 5.1. This measure is intended to favor curves that are relatively close to each other (zeroth order continuity) with a small variation of their end tangents (first order continuity). Two curves are considered for grouping if the energy they produce is less than a given threshold. For each curve, grouping is done with the curve that produces the minimum energy.




Figure 5.1 Grouping using co-curvilinearity

This process is conservative and we find it useful for bridging short breaks. An example of the results obtained in this level is shown in Figure 5.2. In this figure, curve segments are displayed with dots at their extremities. Notice the grouped curve fragments for the cross-section of the cone and the bottom part of the right limb curve of the same object. The break at the top part of the left limb of the vase has also been bridged.




Figure 5.2 Example of results of curve level grouping. a. original contours. b. resulting contours.

hypothesize presence of such cross-sections. A hypothesize-verify process is used to achieve these correspondences.

- Level3: called *SHGC patch level*, where the processed features are local SHGC patches, visible portions, of the surface of an SHGC, delimited by visible parts of surface boundaries (they are defined more rigorously in section 7.1). The motivation for this level is that substantial gaps (due to occlusion or low contrasts) produce incomplete correspondences between surface contours of an SHGC. See Figure 4.2.b for an example. The objective is to form global descriptions of viewed objects in terms of complete SHGC components. For this, another hypothesize-verify process is used.




Figure 4.2 Examples of incomplete correspondences. a. incomplete cross-section. b. incomplete SHGC surface description.

Throughout this hierarchy, the constraints on feature detection, grouping, global verification and completion are based on projective invariant properties of SHGCs. We will discuss the three levels in more detail in the next sections.

5 Curve Level Grouping

The input to this level is an edge image. Using those edges, curves are formed using a linking process based on simple contiguity criteria. Those curves are then segmented at corners. Obtained contours are generally discontinuous. Discontinuities are usually caused by breaks due to low contrasts (where edges are not detectable), errors in localization of edges and occlusion. The gaps between curve segments can be short or large depending on the source of the break. At this level, the objective is to account for relatively short breaks between related curves. For this, we use perceptual grouping based on *co-curvilinearity*. A method similar to that of [Mohan 1989] is used in this

proach to account for the usually sparse features detected in an imperfect contour image. The method assumes that the projection geometry is orthographic and that the viewpoint is general; i.e. small changes in the viewing direction do not change perceptual properties in the image.




Figure 4.1 The three levels of the proposed method

As shown in Figure 4.1, our method is organized as a hierarchy of three perceptual grouping levels handling a hierarchy of three features: curves, parallel symmetries and SHGC patches. The processing is essentially bottom-up starting from an input contour image of a real scene and ending with SHGC descriptions of viewed objects. The three levels are briefly summarized here.

- Level1: called *curve level*, where the processed features are curve segments. The motivation for this level is that contours are often broken. The objective is to form global contours by bridging short gaps using curve grouping based on co-curvilinearity.
- Level2: called *parallel symmetry level*, where the processed features are parallel symmetries between contours. The motivation for this level is that large gaps can separate related contours and that some contours may have been segmented at corners yet they constitute a whole entity (such as cross-sections of GCs). Figure 4.2.a gives an example of a large gap in the cross-section. The objective is to form global correspondences that can be used to

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 $|\underline{V}_2| / |\underline{V}_1| = a$ and (i.e. the ratio of their lengths is *constant* and equal to the *scaling* of the correspondence).

The proof is given in appendix A.3.




Figure 3.3 Projective invariant properties of SHGCs

The usage of these properties will be discussed throughout the description of the method in the next sections.

4 Overview of the Approach

Our method aims at detecting SHGC objects from real image contours. It addresses SHGCs as individual objects, not as part of compound objects. This work, however, lays the foundation of a more general approach that handles compound objects (which we will address in a subsequent effort). The method handles scenes with image imperfections including broken contours, markings and occlusion. It is based on two fundamental aspects. The first one has been discussed in section 3 and is of a geometric nature where projective invariant properties of SHGCs are studied and derived. This is essential since all that is given are 2-D contours of a projected scene. The second one relates to handling the previously mentioned imperfections by using a perceptual grouping ap-

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 [Ulupinar 1990a].

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 section 4 of [Shafer 1983].

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 [Shafer 1983]. In 2-D, tangents to the projections of limbs intersect on the projection of the axis at a common point [Ponce 1989; Ulupinar 1990a].

The properties we add have been reported without proofs in an overview of this work in [Nevatia 1992]. Equivalent ones have been independently derived by [Sato 1992a and b]. Here, we state the new properties and give their proofs.

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*.

The proof of this theorem is given in appendix A.1. This property has proved more robust than property P3 when applied to real image contours. This will be discussed later in section 7.1.

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*.

The proof of this corollary is given in appendix A.2.

is scaled by an amount $r(s)$. For an LSHGC this function is linear; i.e. $r(s) = a (s - s_0)$. Curves of constant t are called *meridians* and curves of constant s are called *cross-sections* (also *parallels*).




Figure 3.1 SHGC representation and terminology

Definition 2: Two planar unit speed curves² $C_1(w_1)$ and $C_2(w_2)$ are said to be *parallel symmetric* [Ulupinar 1990b] if there exists a continuous and monotonic 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 $C_i(w_i)$. Thus, corresponding points have parallel tangent vectors.

The correspondence is said to be linear if f is a linear function. In this case the two curves are similar up to scale and translation. The axis is the locus of midpoints of lines of symmetry (correspondence lines). Figure 3.2. gives an example. A property of linear parallel symmetric curves is that lines of correspondences are either mutually parallel (for a unit scaling) or all intersect at one point (apex). In this work, we only consider linear correspondences for reasons discussed later in this section.




Figure 3.2 Example of parallel symmetry

Now we state projective invariant properties of SHGCs. Those that have been derived in previous work [Ponce 1989, Shafer 1983, Ulupinar 1990a] are stated without proofs. The ones that we introduce here are proved in Appendix A. Figure 3.3 illustrates the properties.

2. a curve is unit speed if it is parameterized by arclength.

cessing modules. The first one detects ends of SHGCs by finding symmetries between contours. The second one detects meridians and the axis. The third module consists of finding the cross-section and the fourth one of recovering 3-D information. Their method and ours are quite similar in the principle of using projective properties. However, they differ in the way those properties are used and in the complexity of the scenes they can handle. Most notably, their system does not handle occlusion. We will compare it to ours in more detail in section 8.

3 Properties of SHGCs

Projective invariant properties have been the focus of research of many researchers, most notably for the problem of *recognition* of objects from their projected contours. Those properties provide strong constraints that, when verified, suggest the presence of corresponding objects in the scene. In the context of shape from contour, projective invariant properties, when they exist, provide strong constraints for the detection of generic shapes. Thus, detection of contours satisfying those constraints is an important step of the figure-ground problem.

SHGCs and their properties have been studied by several researchers in the last few years [Ponce 1989; Shafer 1983; Ulupinar 1990a]. We include the relevant properties from previous work and new properties we have derived in the discussion below. First, we give relevant definitions.

Definition 1: An SHGC (straight homogeneous generalized cylinder) is a surface described by the triplet $\{C, A, S\}$; where C is called the *cross-section* function representing a planar curve, A is called the *axis* representing a straight line in 3-D (not necessarily orthogonal to the cross-section plane) and S is called the *scaling* function, giving for each point on the axis the amount of scaling of the curve C .

In casual terms, an SHGC is the surface obtained by sweeping C along A while transforming it (scaling it in this case) by S . Let $C(t) = (u(t), v(t))$ be a parametrization of C , $r(s)$ the scaling function and α the angle between the cross-section plane and the SHGC axis (s -direction), then the surface of the SHGC can be parameterized as follows (using the formulation of [Shafer 1983]):

$$S(t, s) = (u(t) r(s) \sin\alpha, v(t) r(s), s + u(t) r(s) \cos\alpha) \quad (3.1)$$

When $\alpha = \pi / 2$, we obtain a *right* SHGC (RSHGC). Figure 3.1 shows an RSHGC and the chosen configuration of the axes. Equation (3.1) implies that for each value of s , the cross-section

Few efforts have been made in the research community for the detection of generic objects from real 2-D image contours. The objects addressed include GCs and ribbons (2-D counterpart of GCs). We briefly discuss some of the proposed methods.

ACRONYM [Brooks 1983] was one of the first such methods. In this system, stored models are used for the interpretation of image contours. Volumetric models based on GCs are used along with partially ordered non-linear algebraic constraints on model parameters in order guide the interpretation process. Successful interpretation (and recognition) of object contours is achieved through a constraint-based matching process between predicted image and model features. This system has been applied to straight GCs (airplane models) with a top view and is limited to the stored models available.

Ponce et. al [1989] have proposed a method to detect the axis of an SHGC from its contour. The method is based on a projective invariant property they derived that relates corresponding points of an SHGC limb contours. The property is that tangents to limb projections at such points intersect along a straight line which is the projection of the SHGC axis. Using this property and another one which relates zero-curvature points, they proposed two methods for the detection of the projection of the axis of an SHGC. The first algorithm uses a Hough transform and the second one requires that the sweeping function have at least two inflection points. Robustness of that method to noisy contours and occlusion is not discussed. It is to be expected, however, that localization of zero-curvature points is sensitive to contour distortions.

Rao and Nevatia [1989] and Mohan and Nevatia [1989] have proposed two different method for the figure-ground problem in complex scenes with occlusion, in the context of ribbons. Those methods use rather intuitive definitions of symmetries and grouping constraints. For example, in the latter method, symmetries necessarily join curve extremities. Thus different correspondences are obtained for different occlusion patterns. Both methods address the problem of selecting ribbons and symmetries, a necessary task for the figure-ground problem. The former method uses a graphical representation of detected ribbons followed by an analysis of cycles of that graph subject to some regularity constraints. The latter one uses a Hopfield network to select the ‘best’ symmetries relating curves bounding ribbon surfaces. Both methods have produced impressive results on complex scenes but, as mentioned above, their constraints are intuitive. Their results can, therefore, hardly be used to generate the necessary (rigorous) constraints needed for 3-D shape recovery.

Sato and Binford [1992a, 1992b] have recently proposed a method to detect SHGCs. Their method also exploits projective properties of SHGCs for their detection. For this, they use four pro-

erties in a bottom up, perceptual grouping approach handling a hierarchy of three levels of features: curves, symmetries and SHGC surface patches.

The input to the method is a real contour image and the output SHGC descriptions of viewed objects. In the first level, the objective is to form global contours. Grouping of contour fragments based on co-curvilinearity is used for that purpose. In the second level, the objective is to form global symmetry correspondences between contours. The aim is to hypothesize presence of cross-sections. The third level addresses SHGC surface patch features. The objective is to form complete SHGC descriptions from sparse, visible local surface patches. Throughout this hierarchy, relevant features are detected and missing ones are inferred. Constraints on local feature detection, feature grouping, global consistency and feature completion are based on projective invariant properties of SHGCs.

We have implemented a system that produced satisfactory results on rather complex scenes by the standards of currently developed methods. These results can be directly used for 3-D shape recovery (by the method of [Ulupinar 1990a], for example). Thus, our method complements this latter one by providing the necessary 2-D descriptions required as input. The obtained descriptions also have potential applications in recognition using structured descriptions of scene objects in terms of GC components and their relationships.

The remainder of this paper is organized as follows. In section 2, we discuss related previous work. In section 3, we discuss the projective invariant properties of SHGCs we use. In section 4, we introduce in more detail the organization of our approach. Sections 5, 6 and 7 discuss in detail the different levels of our hierarchy. Examples of obtained results will also be given. In section 8, we discuss the capabilities and limitations of our method and compare them with those of other related methods. In section 9, we demonstrate the usage of those results for 3-D shape recovery. We conclude this paper in section 10.

2 Previous Work

Related work to the problem of handling imperfect contours includes methods of perceptual grouping and methods of generic shape detection. These latter two are not disjoint. For lack of space, we will limit the discussion to the more related work on generic shape detection. Some of the discussed methods do use perceptual grouping.

As discussed previously, methods of shape from contour handle a given class of objects (shape model). In order to address the problem of shape from imperfect contours, an adequate shape model is needed. Such a model should be general enough so that it covers a broad class of objects, yet have well defined properties so that their detection, description and subsequent 3-D shape recovery and recognition can be done in a rigorous way. A number of such *generic models* have been used in the field for shape description.

The generalized cylinder (GC) is one such generic model. Since their introduction by Binford [1971], GCs have been studied and used by many researchers. The advantages of using GCs are numerous. First, they satisfy most requirements of a ‘good’ shape description scheme by providing rich, stable¹ and structured descriptions. A large class of objects we see daily are made up of GC components. The types of such components are usually drawn from a small set of primitives having different attributes (axis shape for example). Thus, it is sufficient to address a small set of primitive GCs for an approach that addresses a large number of complex objects. Second, successful results have already been obtained for 3-D shape recovery from contours for certain classes of GCs [Ulupinar 1990a, 1990b, 1991b, 1992; Zerroug 1993b]. Finally, a number of researchers have derived important projective invariant properties of GCs [Ponce 1989; Rao 1988; Shafer 1983; Ulupinar 1991a]. Those invariants have been successfully used for 3-D shape recovery in the latter reference. However, little work has been done on the detection of generic objects from real image contours. The few methods differ in the class of objects they address and in the complexity of the scenes they can handle. We discuss some of them in section 2. None of those methods, however, handles GCs in complex scenes with occlusion.

In this work, we address the problem of shape description and scene segmentation (figure-ground problem) in the presence of broken contours, markings and occlusion for an important subclass of GCs: SHGCs. These latter are defined as the surface obtained by sweeping a planar cross-section along a straight axis while scaling it (a more formal definition is given in section 3). SHGCs are one class of primitive GCs that belong to the primitive set mentioned above. Our approach exploits *projective invariant properties* of SHGCs in order to guide the segmentation and description processes. We believe that invariant properties of generic shapes greatly help solving the figure-ground problem. In fact, an essential characteristic of the segmentation of an image of a 3-D scene should be its viewpoint invariance. The method we propose consists of using those rigorous prop-

1. with respect to changes in viewing conditions, including viewpoint.

produce contours. Figure 1.1 shows an example of a real image and its extracted edges by a Canny edge detector [Canny 1986]. The difficulty in dealing with such imperfections is that it is impossible to tell, by just looking at the contours individually, which constitute real contours and of what objects and which do not; i.e. objects versus background. This problem is known as the *figure-ground* problem. To address this problem, it is necessary to use some form of grouping process that groups relevant features together, but also discards irrelevant ones.




Figure 1.1 A real image and its extracted edges

Among the few efforts in this direction are works based on perceptual grouping, such as [Dolan 1989; Hummel 1992; Lowe 1985; Mohan 1989]. Most of those methods address curve features. Simple compatibility constraints such as co-curvilinearity and collinearity have proved quite useful for grouping broken but related contours. However, local constraints between curves are insufficient (sometimes wrong) to generate all the relevant groupings, as perceived features are not only curves but (global) relationships, such as symmetries, between curves as well. Given the hierarchical nature of object and scene features, it is important to address the figure ground problem by using feature grouping at different levels of this hierarchy. In doing so, the interaction between local and global features causes segmentation and description to be done in a more principled way. In order to be of interest for imperfect contours, the grouping constraints should be locally applicable so as to handle occlusion and gaps. Yet, they should provide global criteria that discriminate between true instances and accidental ones. In this paper, we will describe our approach to the figure ground problem based on these observations.

1 Introduction

One of the fundamental problems in computer vision is the recovery of the shape of the objects in a scene. Monocular intensity images are one of the most easily available and convenient sources of data for such a task. However, detection of shape from a monocular intensity image is particularly difficult as there is ambiguity caused by projection. Human vision does show that substantial information can be inferred from a single intensity image. This includes both segmentation of the scene into different objects and perception of their 3-D shape. To achieve such ability in machines has continued to be one of the most challenging problems in computer vision.

Many cues are available for inferring 3-D shape from 2-D images. They include contours, shading and texture. We believe that contour is by far the most dominant and reliable of them. This can be justified on a psychological basis by human experiments [Biederman 1987] as well as on computational grounds. All shape from 2-D image methods require some assumptions. We believe that those using contour are much less restrictive than those using shading, for example, where complete reflectance functions need to be known and assumed to be constant over the scene.

Work on shape from contour has started since the early days of computer vision. Early efforts have addressed polyhedral objects through such works as [Clowes 1971; Mackworth 1973; Kanade 1981]. Subsequent efforts, such as [Barrow 1981; Binford 1971; Brooks 1983; Gross 1990; Horaud 1988; Malik 1987; Nalwa 1989; Richetin 1991; Sato 1992b; Ulupinar 1990a, 1990b, 1991b, 1992; Zerroug 1993b], have addressed curved objects. These latter introduce more difficulties as some of their contours, such as limbs and cusps, are viewpoint dependent. Most of the proposed methods consist of using properties of the observed contours to generate constraints on 3-D shape. Such properties are usually based on knowledge about the shape model used (type of shapes handled), differential geometry of curved surfaces and the projection geometry. An example is the approach of Ulupinar and Nevatia [1990a, 1990b, 1991b, 1992] which handles objects such as Zero Gaussian Curvature (ZGC) surfaces, SHGCs (straight homogeneous generalized cylinders), PRCGCs (planar right constant generalized cylinders) and multiple ZGC-surface objects.

However, most previous work on inferring 3-D shape from 2-D assumes that the problem of object (surface) segmentation has been solved, whereas this is perhaps the most critical step in monocular scene analysis. The difficulty lies in the interpretation of real images. These latter usually produce contours with many imperfections such as noise, distortions, breaks and occlusion. Further, not just 'real' object contours are present in an image. Surface markings and shadows also