Recognising rotationally symmetric surfaces from their outlines*

David A. Forsyth[1], *Joseph L. Mundy*[2], *Andrew Zisserman*[3] *and Charles A. Rothwell*[3]

[1] Department of Computer Science, University of Iowa, Iowa City, Iowa, USA.
[2] The General Electric Corporate Research and Development Laboratory, Schenectady, NY, USA.
[3] Robotics Research Group, Department of Engineering Science, Oxford University, England.

Abstract. Recognising a curved surface from its outline in a single view is a major open problem in computer vision. This paper shows techniques for recognising a significant class of surfaces from a single perspective view. The approach uses geometrical facts about bitangencies, creases, and inflections to compute descriptions of the surface's shape from its image outline. These descriptions are unaffected by the viewpoint or the camera parameters. We show, using images of real scenes, that these representations identify surfaces from their outline alone. This leads to fast and effective recognition of curved surfaces.

The techniques we describe work for surfaces that have a rotational symmetry, or are projectively equivalent to a surface with a rotational symmetry, and can be extended to an even larger class of surfaces. All the results in this paper are for the case of full perspective. The results additionally yield techniques for identifying the line in the image plane corresponding to the axis of a rotationally symmetric surface, and for telling whether a surface is rotationally symmetric or not from its outline alone.

1 Introduction

There has been a history of interest in recognising curved surfaces from their outlines. Freeman and Shapira [7], and later Malik [9] investigated extending line labelling to curved outlines. Brooks [2] studied using constraint-based modelling techniques to recognise generalised cylinders. Koenderink [8] has pioneered work on the ways in which the topology of a surface's outline changes as it is viewed from different directions, and has studied the way in which the curvature of a surface affects the curvature of its outline. Ponce [12] studied the relationships between sections of contour in the image of a straight homogenous generalised cylinder. Dhome [4] studied recognising rotationally invariant objects by computing pose from the image of their ending contours.

Terzopolous *et al.* [17] compute three-dimensional surface approximations from image data, based around a symmetry seeking model which implicitly assumes that "the axis of the object is not severely inclined away from the image plane" (p. 119). These approximations can not, as a result, be used for recognition when perspective effects are

* DAF acknowledges the support of Magdalen College, Oxford, of the University of Iowa, and of GE. JLM acknowledges the support of the GE Coolidge Fellowship. AZ acknowledges the support of SERC. CAR acknowledges the support of GE. The GE CRD laboratory is supported in part by the following: DARPA contract DACA-76-86-C-007, AFOSR contract F49620-89-C-003.

significant. Despite this work, it has been hard to produce robust, working model based vision systems for curved surfaces.

Ponce and Kriegman [13] show that elimination theory can be used to predict, in symbolic form, the outline of an algebraic surface viewed from an arbitrary viewing position. For a given surface a viewing position is then chosen using an iterative technique, to give a curve most like that observed. The object is then recognized by searching a database, and selecting the member that gives the best fit to the observed outline. This work shows that outline curves strongly constrain the viewed surface, but has the disadvantage that it cannot recover surface parameters without solving a large optimization problem, so that for a big model base, each model may have to be tested in turn against the image outline.

A number of recent papers have shown how *indexing functions* can be used to avoid searching a model base (e.g. [5, 18, 14]). Indexing functions are descriptions of an object that are unaffected by the position and intrinsic parameters of the camera, and are usually constructed using the techniques of invariant theory. As a result, these functions have the same value for any view of a given object, and so can be used to index into a model base without search. Indexing functions and systems that use indexing, are extensively described in [10, 11], and [16] displays the general architecture used in such systems.

To date, indexing functions have been demonstrated only for plane and polyhedral objects. Constructing indexing functions for curved surfaces is more challenging, because the indexing function must compute a description of the surface's shape from a single outline. It is clear that it is impossible to recover global measures of surface shape from a single outline if the surfaces involved are unrestricted. For example, we can disturb any such measure by adding a bump to the side of the surface that is hidden from the viewer. An important and unresolved question is how little structure is required for points to yield indexing functions.

In this paper, we emphasize the structure of image points by demonstrating useful indexing functions for surfaces which have a rotational symmetry, or are within a 3D projectivity of a surface with a rotational symmetry. This is a large and useful class of surfaces.

2 Recognising rotationally symmetric surfaces

In this section, we show that lines bitangent to an image contour yield a set of indexing functions for the surface, when the surface is either rotationally symmetric, or projectively equivalent to a rotationally symmetric surface. This follows from a study of the properties of the outline in a perspective image.

2.1 Geometrical properties of the outline

The outline of a surface in an image is given by a system of rays through the camera focal point that are tangent to the surface. The points of tangency of these rays with the surface form a space curve, called the *contour generator*. The geometry is illustrated in figure 1.

Points on the contour generator are distinguished, because the plane tangent to the surface at such points passes through the focal point (this is an alternative definition of the contour generator). As a result, we have:

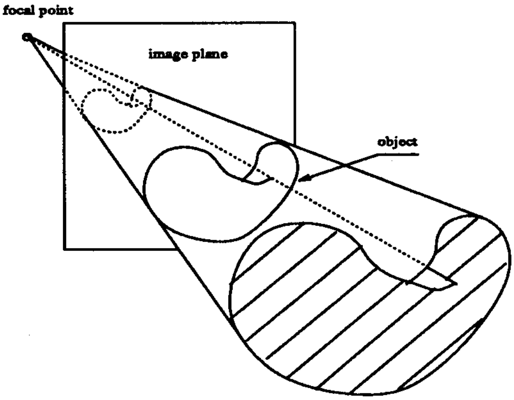

Fig. 1. The cone of rays, through the focal point and tangent to the object surface, that forms the image outline, shown for a simple object.

Lemma: Except where the image outline cusps[4], a plane tangent to the surface at a point on the contour generator (by definition, such a plane passes through the focal point), projects to a line tangent to the surface outline, and conversely, a line tangent to the outline is the image of a plane tangent to the surface at the corresponding point on the contour generator.

As a corollary, we have:

Corollary 1: A line tangent to the outline at two distinct points is the image of a plane through the focal point and tangent to the surface at two distinct points, both on the contour generator.

This yields useful relationships between outline properties and surface properties. For example:

Corollary 2: The intersection of two lines, bitangent to the outline is a point, which is the image of the intersection of the two bitangent planes represented by the lines[5].

The lemma and both corollaries follows immediately from considering figure 2. Generic surfaces admit one-parameter systems of bitangent planes, so we can expect to observe and exploit intersections between these planes.

One case in which the intersections are directly informative occurs when the surface is rotationally symmetric. The envelope of the bitangent planes must be either a right circular cone, or a cylinder with circular cross-section (this is a right circular cone whose vertex happens to be at infinity). We shall draw no distinction between vertices at infinity and more accessible vertices, and refer to these envelopes as *bitangent cones.* These comments lead to the following

[4] We ignore cusps in the image outline in what follows.

[5] **Proof:** Each of the bitangent planes passes through the focal point, so their intersection must pass through the focal point, and in particular is the line from the focal point through the intersection of the bitangent lines.

Key result: The vertices of these bitangent cones must lie on the axis (by symmetry), and so are collinear. Assuming the focal point lies outside the surface, as figure 2 shows, the vertices of the bitangent cones ***can be observed in an image.*** The vertices appear as the intersection of a pair of lines bitangent to the outline.

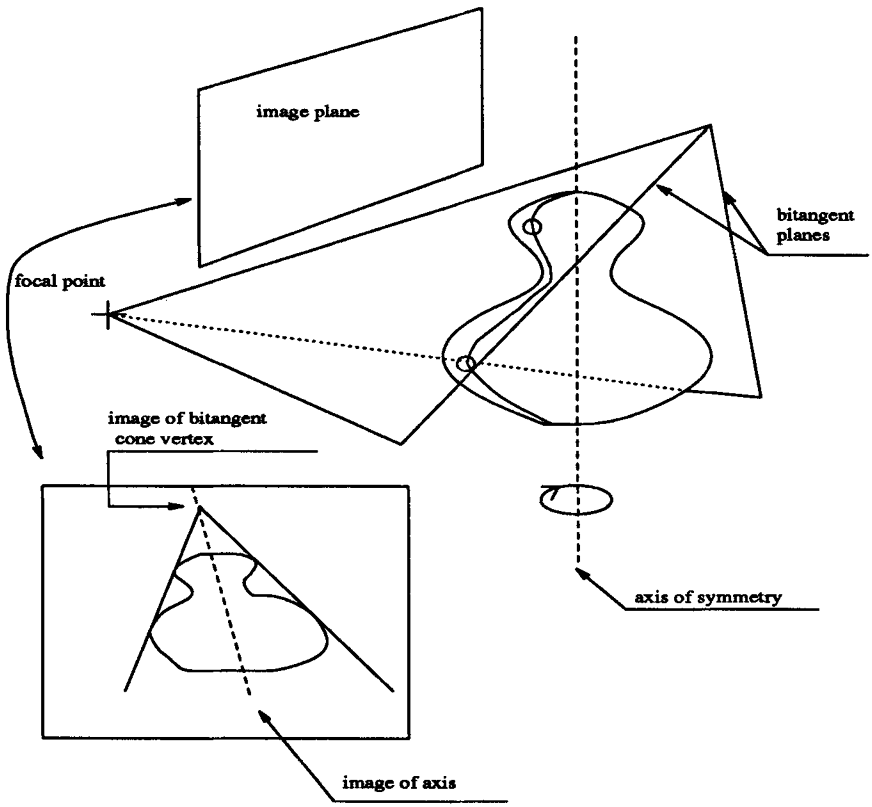

Fig. 2. A rotationally symmetric object, and the planes bitangent to the object and passing through the focal point, are shown. It is clear from the figure that the intersection of these planes is a line, also passing through the focal point. Each plane appears as a line in the image: the intersection of the planes appears as a point, which is the image of the vertex of the bitangent cone. Note in particular that the image outline *has no symmetry.* This is the generic case.

As a result, if the surface has four or more bitangent cones, the vertices yield a system of four or more collinear points, lying on the axis of the surface. These points project to points that are collinear, and lie on the image of the axis of symmetry. These points *can be measured in the image.* This fact yields two important applications:

- Cross-ratios of the image points, defined below, yield indexing functions for the surface, which can be determined from the outline alone.
- The image points can be used to construct the image of the axis of a rotationally symmetric surface from its outline.

The second point can be used to extend work such as that of Brady and Asada [1] on symmetries of frontally viewed plane curves to considering surface outlines.

We concentrate on the first point in this paper. The map taking these points to their corresponding image points is a projection of the line, and so the projective invariants of a system of points yield indexing functions. A set of four collinear points A, B, C, D has a projective invariant known as its *cross ratio*, given by:

$$\frac{(AC)(BD)}{(AD)(BC)}$$

where AB denotes the linear distance from A to B. The cross ratio is well known to be invariant to projection, and is discussed in greater detail in [10]. The cross ratio depends on the order in which the points are labeled. If the labels of the four points are permuted, a different value of the cross ratio results. Of the 24 different labeling possibilities, only 6 yield distinct values. A symmetric function of cross ratios, known as a j-invariant is invariant to the permutation of its arguments as well as to projections [11]. Since a change in camera parameters simply changes the details of the projection of the points, but does not change the fact that the map is a projection, these cross-ratios are invariant to changes in the camera parameters.

A further result follows from the symmetries in figure 2. It is possible to show that, although the outline does not, in general, have a symmetry, it splits into two components, which are within a plane projectivity of one another. This means that the techniques for computing projective invariants of general plane curves described in [15], can be used to group corresponding outline segments within an image.

2.2 Indexing functions from images of real objects

We demonstrate that the values of the indexing functions we have described are stable under a change in viewing position, and are different for different objects. In figure 3, we show two images each of two lampstands, taken from different viewpoints. These images, demonstrate that the outline of a rotationally symmetric object can be substantially affected by a change in viewpoint. Three images of each lampstand were taken in total, including these images. For each series of images of each lampstand, bitangents were constructed by hand, and the bitangents are shown overlayed on the images. The graph in figure 4 shows the cross-ratios computed from the vertices in each of three images of each of two lampstands. The values of the cross ratio are computed for only one ordering of the points, to prevent confusion. As predicted in [5], the variance of the larger cross-ratio is larger; this effect is discussed in [5], and is caused by the way the measurements are combined in the cross-ratio. The results are easily good enough to distinguish between the lampstands, from their outlines alone.

3 Generalizing the approach

This approach can be generalized in two ways. Firstly, there are other sources of vertices than bitangent lines. Secondly, the geometrical construction described works for a wider range of surfaces than the rotationally symmetric surfaces. We will demonstrate a range of other sources of vertices assuming that the surface is rotationally symmetric, and then generalize all the constructions to a wider range of surfaces in one step.

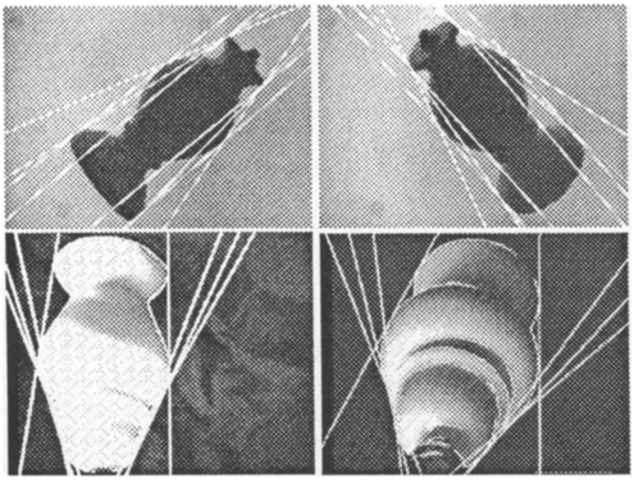

Fig. 3. This figure shows two views each of two different lampstands. Bitangents, computed by hand from the outlines, are overlaid.

3.1 Other sources of vertices

Other sources of vertices, illustrated in figure 5, are:

- **The tangents at a crease or an ending in the outline:** We assume that we can distinguish between a crease in the outline, which arises from a crease in the surface, and a double point of outline, which is a generic event that may look like a crease. In this case, these tangents are the projections of planes tangent to the surface, at a crease in the surface.
- **A tangent that passes through an ending in the outline:** These are projections of planes that are tangent to the surface, and pass through an ending in the surface.
- **Inflections of the outline:** These are projections of planes which have three-point contact with the surface.

In each case, there is a clear relationship between the tangent to the outline and a plane tangent to the surface, and the envelope of the system of planes tangent to the surface and having the required property, is a cone with a vertex along the axis. These sources of information are demonstrated in figure 5. These results can be established by a simple modification of the argument used in the bitangent case.

3.2 Generalizing to a wider class of surfaces

We have constructed families of planes tangent to a rotationally symmetric surface and distinguished by some property. The envelope of each family is a cone, whose vertex lies on the axis of the surface. The projections of these vertices can be measured in the image, by looking at lines tangent to the outline. To generalize the class of surfaces to which these constructions apply, we need to consider the properties we are using, and how they behave under transformation. The properties used to identify vertices are preserved under projective mappings of space. By this we mean that, using bitangency as an example, if we take a surface and a bitangent plane, and apply a projectivity of space to each, the new plane is still bitangent to the new surface, and the old points of tangency will map

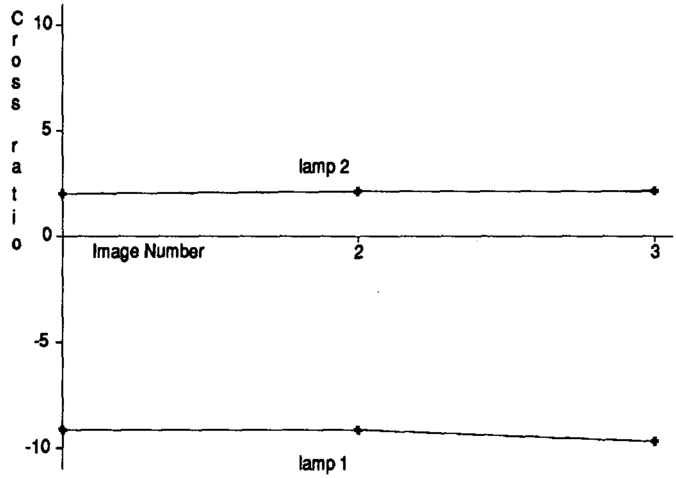

Fig. 4. A graph showing one value of the cross-ratio of the vertex points for three different images, taken from differing viewpoints, each of two different vases. This figure clearly shows that the values of the indexing functions computed are stable under change of view, and change for different objects, and so are useful descriptors of shape. As expected (from the discussion in [5]), the variance in a measurement of the cross ratio increases as its absolute value increases.

to the new points of tangency. The other properties are preserved because projectivities preserve incidence and multiplicity of contact.

These results mean that if we take a rotationally symmetric surface, and apply a projective mapping, the cones and vertices we obtain from the new surface are just the images of the cones and vertices constructed from the old surface. Since a projective mapping takes a set of collinear points to another set of collinear points, we can still construct indexing functions from these points. This means that, for our constructions to work, the surface need only be projectively equivalent to a rotationally symmetric surface. One example of such a surface would be obtained by squashing a rotationally symmetric surface so that its cross section was an ellipse, rather than a circle. This result substantially increases the class of surfaces for which we have indexing functions that can be determined from image information.

A further generalisation is possible. The cross-ratio is a remarkable invariant that applies to sets of points lying on a wide range of algebraic curves. If a curve supports a one-to-one parametrisation using rational functions, a cross-ratio can be computed for a set of four points on that curve. This follows because one can use the parametrisation to compute those points on the line that map to the points distinguished on the curve, and then take the cross-ratio of the points on the line[6]. Curves that can be parametrised are also known as curves with genus zero. There is a wide range of such curves; some examples in the plane include a plane conic, a cubic with one double point and a quartic with either one triple point or three double points. In space, examples include the twisted

[6] The parametrisation is birational. Any change of parametrisation is a birational mapping between lines, hence a projectivity, and so the cross-ratio is well defined.

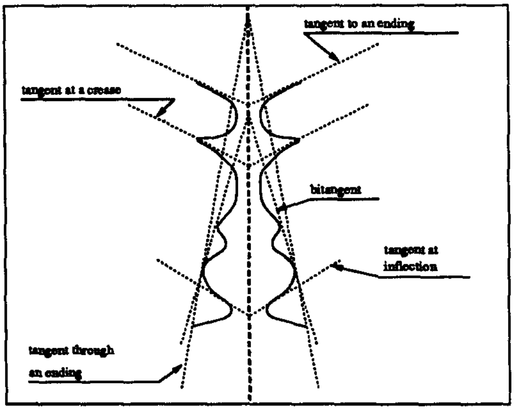

Fig. 5. The known cases that produce usable coaxial vertices. Note that although this figure appears to have a reflectional symmetry, this is not a generic property of the outline of a rotationally symmetric object.

cubic, and curves of the form $(t, p(t), q(t), 1)$, where p and q are polynomials (the points at infinity are easily supplied).

Remarkably, the resulting cross-ratio is invariant to projectivities and to projection from space onto the plane (in fact, to any birational mapping). This means that, for example, if we were to construct a surface for which the bitangent vertices or other similarly identifiable points lie on such a curve, that surface could easily be recognised from its outline in a single image, because the cross-ratio is defined, and is preserved by projection. Recognition would proceed by identifying the image of the bitangent vertices, identifying the image of the projected curve that passes through these points, and computing the cross-ratio of these points on that curve, *which would be an invariant.*

Since there is a rich range of curves with genus zero, this offers real promise as a modelling technique which has the specific intent of producing surface models that are both convincing models of an interesting range of objects, and intrinsically easy to recognise.

4 Conclusions

We have constructed indexing functions, which rely only on image information, for a useful class of curved surfaces. We have shown these functions to be useful in identifying curved objects in perspective images of real scenes.

This work has further ramifications. It is possible to use these techniques to determine whether an outline is the outline of a rotationally symmetric object, and to determine the image of the axis of the object. As a result, it is possible to take existing investigations of the symmetry properties of image outlines, and extend them to consider surface properties, measured in a single image, in a principled way.

As we have shown, image information has deep geometric structure that can be exploited for recognition. In fact, image outlines are so rich that recent work at Iowa[6] has shown that a generic algebraic surface of any degree can be recovered up to a projective mapping, from its outline in a single image.

References

1. Brady, J.M. and Asada, H., "Smoothed Local Symmetries and their implementation," *IJRR-3*, 3, 1984.
2. Brooks, R. A., "Model-Based Three-Dimensional Interpretations of Two Dimensional Images," IEEE PAMI, Vol. 5, No. 2, p. 140, 1983.
3. Canny J.F. "Finding Edges and Lines in Images," TR 720, MIT AI Lab, 1983.
4. Dhome, M., LaPreste, J.T, Rives, G., and Richetin, M. "Spatial localisation of modelled objects in monocular perspective vision," *Proc. First European Conference on Computer Vision*, 1990.
5. D.A. Forsyth, J.L. Mundy, A.P. Zisserman, A. Heller, C. Coehlo and C.A. Rothwell (1991), "Invariant Descriptors for 3D Recognition and Pose," *IEEE Trans. Patt. Anal. and Mach. Intelligence*, **13**, 10.
6. Forsyth, D.A., "Recognising an algebraic surface by its outline," Technical report, University of Iowa Department of Computer Science, 1992.
7. H. Freeman and R. Shapira, "Computer Recognition of Bodies Bounded by Quadric Surfaces from a set of Imperfect Projections," IEEE Trans. Computers, **C27**, 9, 819-854, 1978.
8. Koenderink, J.J. *Solid Shape*, MIT Press, 1990.
9. Malik, J., "Interpreting line drawings of curved objects," *IJCV*, **1**, 1987.
10. J.L. Mundy and A.P. Zisserman, "Introduction," in J.L. Mundy and A.P. Zisserman (ed.s) *Geometric Invariance in Computer Vision*, MIT Press, 1992.
11. J.L. Mundy and A.P. Zisserman, "Appendix" in J.L. Mundy and A.P. Zisserman (ed.s) *Geometric Invariance in Computer Vision*, MIT Press, 1992.
12. Ponce, J. "Invariant properties of straight homogenous generalised cylinders," *IEEE Trans. Patt. Anal. Mach. Intelligence*, **11**, 9, 951-965, 1989.
13. J. Ponce and D.J. Kriegman (1989), "On recognising and positioning curved 3 dimensional objects from image contours," *Proc. DARPA IU workshop*, pp. 461-470.
14. Rothwell, C.A., Zisserman, A.P., Forsyth, D.A. and Mundy, J.L., "Using Projective Invariants for constant time library indexing in model based vision," *Proc. British Machine Vision Conference* , 1991.
15. Rothwell, C.A., Zisserman, A.P., Forsyth, D.A. and Mundy, J.L., "Canonical frames for planar object recognition," *Proc. 2nd European Conference on Computer Vision*, Springer Lecture Notes in Computer Science, 1992.
16. Rothwell, C.A., Zisserman, A.P., Forsyth, D.A. and Mundy, J.L., "Fast Recognition using Algebraic Invariants," in J.L. Mundy and A.P. Zisserman (ed.s) *Geometric Invariance in Computer Vision*, MIT Press, 1992.
17. Terzopolous, D., Witkin, A. and Kass, M. "Constraints on Deformable Models: Recovering 3D Shape and Nonrigid Motion," *Artificial Intelligence*, **36**, 91-123, 1988.
18. Wayner, P.C. "Efficiently Using Invariant Theory for Model-based Matching," Proceedings CVPR, p.473-478, 1991.

This article was processed using the LTEX macro package with ECCV92 style