Using Occluding Contours for 3D Object Modeling*

Régis Vaillant
INRIA
Sophia-Antipolis 2004 Route des Lucioles
06565 Valbonne Cedex France

Abstract

This paper presents an algorithm for detecting the occluding contours generated by a surface and for reconstructing depth along them. It also describes an algorithm for computing the two main curvatures of the surface in the neighborhood of the occluding contours. We have used these algorithms on synthetic and real data.

1 Introduction

One of the aims of computer vision is to extract concise surface descriptions from several images of a scene. The descriptions can be used for the purpose of object recognition and for geometric reasoning (such as obstacle avoidance).

Stereovision is often used for recovering the structure of the 3D world. Standard techniques can determine the depth of edges on a surface. These techniques fail with extremal boundaries as these change according to the viewpoint.

Nonetheless, in some cases, they are the only source of 3D information (imagine a white sphere on a black background), if we are not willing to exploit shape from shading techniques. In all cases, they are a rich source of 3D information as will be shown here. In this paper, we propose a new method for detecting extremal boundaries. We also propose an algorithm for reconstructing exactly the curves observed by each camera and computing the principle curvatures of the object surface in their vicinity.

In the first part, we briefly describe the main characteristics of the experimental setup used and we present the theoretical framework of our algorithms. In the second part, we present a method for detecting and reconstructing the extremal boundaries. The third part is devoted to the study of the computation of the two first fundamental forms of the surface in the neighbourhood of the extremal boundary. In the last part, we present results on real and synthetic data and a discuss their accuracy.

*This work was supported in part by Esprit project P2502,`Voila`

2 Background

We assume that after calibration our cameras can be accurately modeled as pinholes. We suppose that we are looking at a smooth object, i.e., whose surface is at least C^2. For a given position of the camera, we can draw the optical rays tangent to the surface of the object. These rays cut on the retinal plane a curve, the *occluding contour*, and touch the object along a smooth curve on its surface, the *rim*.

Several questions can be asked at this point. First what kind of information can be obtained from one occluding contour and second what kind of information can be obtained from several, possibly many, occluding contours obtained from a number of different viewpoints. The first question has been dealt with by Koenderink [5]. In his paper, he proves that concavities and convexities of the visual contour allow to draw implications about the local shape of the surface looked at: convexity of the contour corresponds to a convex patch of the surface while a concavity correspond to a saddle-shaped patch. These conclusions fall from a nice theorem which has also been derived later by Brady [1].

The second question has been addressed by a number of authors, among which Giblin [3], who worked theoretically and Basri and Ullman [6] who worked on the positioning of objects from their occluding contours. Giblin proposes to consider the surface of the object as the envelope of its tangent planes. There are two problems with this: how to compute the envelope of a family of planes and how to handle inflection points. Giblin and Weiss propose to solve the problem by assuming a planar motion of the camera. In a second approach, they derive information about the surface from singular points of the occluding contours.

We also consider the surface of the object as the envelope of its tangent planes but make no assumption about the camera motion or about the projection on the retina plane being orthographic. In fact we deal with the full perspective projection case.

2.0.1 Definitions and notations:

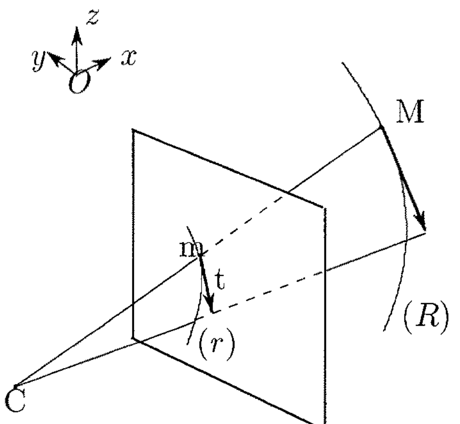

Figure 1: A Rim (R) and its image (r)

As shown in figure 1, we consider a fixed coordinate system $(Oxyz)$; the optical center is at C. The camera looks at the rim (R) which produces the occluding contour (r). A point m on (r) is the image of a point M on (R) at which the optical ray determined by Cm is tangent to the object surface. The tangent plane to the surface at M is defined by the optical ray and the tangent t to the occluding contour at m. Let n be the unit length normal vector to this plane, defined by its Euler angles θ and ϕ and $p(\theta, \phi)$ the distance from the origin to the tangent plane. The equation of this plane can be written as:

$$n(\theta, \phi)^T X - p(\theta, \phi) = 0 \tag{1}$$

where X is the vector $(x, y, z)^T$ and $n = (\cos(\theta)\cos(\phi), \sin(\theta)\cos(\phi), \sin(\phi))$.

2.0.2 The envelope theorem

Now consider the mapping

$$(\theta, \phi) \rightarrow p(\theta, \phi)$$

which associates to every direction the distance from the origin to the plane tangent to the surface whose normal is in the direction (θ, ϕ). In fact, we know that this mapping is locally one to one for elliptic and hyperbolic points [4] but not for parabolic points.

The envelope of the two parameters family of planes defined by equation (1) is obtained by eliminating θ and ϕ between equation (1) and

$$\frac{\partial n(\theta, \phi)}{\partial \theta}^T X - \frac{\partial p(\theta, \phi)}{\partial \theta} = 0 \tag{2}$$

$$\frac{\partial n(\theta, \phi)}{\partial \phi}^T X - \frac{\partial p(\theta, \phi)}{\partial \phi} = 0 \tag{3}$$

The physical interpretation of this is that the point M where the plane of equation (1) is tangent to the surface is obtained as the intersection with the planes defined by equations (2), and (3).

Mathematically, there are no difficulties; it is from the practical standpoint that they arise. Indeed, in practice we measure pieces of the surface $(\mathcal{P})$ from which we have to estimate first and second order derivatives which in turn yield properties of the object surface.

3 Detection of the occluding contours

In the previous part, we have assumed that we can detect the extremal boundaries. In fact this is not an easy problem. We will show in this part that sophisticated models are needed. This investigation is interesting as it provides us with some ideas about the numerical stability of the algorithms that we want to implement. We use an algorithm which is a simplified version of the general algorithm which will be detailled in the following part. This simplified algorithm, allows us test the feasability of this kind of computations.

3.1 Edge classification

The most interesting property of occluding edges is that they do not correspond to a physical marking on the surface. This means that they do not correspond to a discontinuity of the normal of the surface or in the reflectance properties of the surface. In spite of this, they are detected, as the other edges by the edge detection process.

We suppose that we have matched segments among different images. We want to verify if they belong to an extremal boundary. One way to proceed is to assume that they belong to one and to write the corresponding equations. We make the hypothesis that the observed surface is part of a cylinder. This provides us, with a number of equations that can be used to compute the parameters of the hypothetical cylinder: its axis and its radius. Fortunately this computation can be divided into two independent parts:

- the direction of the axis.

- the position of the axis and the radius.

▷ **Computation of the direction of the axis of the cylinder**

We know the optical plane corresponding to an image line. The axis of the cylinder is solution of a linear equation which is a function of the normal to the optical plane. The problem turns to be equivalent to finding the smallest eigen-vector for a symmetric matrix.

▷ **Computation of the position of the axis and the radius of the cylinder**

These computations are very simple if we perform them in the right coordinate system. A good one is $(Ouvw)$, where w is the direction of the axis of the cylinder, u and v define an arbitrary frame in the plane $\mathcal{P}$ which is perpendicular to w. The projection of the cylinder is a circle $\mathcal{C}$ and the parameters of this cylinder can be obtained by solving linear equations.

We have a set of equations which can be used to compute the parameters of a cylinder such that the observed line segments are the image of its rim, as seen from each camera. We need a criterion to check whether our hypothesis is correct i.e. do we observe the rim of something which is locally cylindrical or a normal edge. We can first notice that the model we used is still correct if we suppose that the radius of the cylinder is zero. A cylinder of zero radius is physically equivalent to a normal edge. The occluding edges and the normal edges can be classified by performing a test on the value of the radius. There is still a problem: we have to fix a threshold for taking a decision.

We want to estimate the uncertainty on the measure of the radius of the cylinder. We can consider that we have constructed a function f such that

$$(c_1, c_2, r) = f(u_1, v_1, \cdots, u_n, v_n)$$

where (u_i, v_i) are the coordinates of the extremities of the image-segments. We suppose that these values are corrupted by a Gaussian noise of variance σ_{u_i} and σ_{v_i}. In this case, we can express the uncertainty on r by the formula:

$$\sigma = \sqrt{\sum_{u_i} \frac{\partial f(u_i, v_i, \ldots)}{\partial u_i}^2 \sigma_{u_i}^2 + \sum_{v_i} \frac{\partial f(u_i, v_i, \ldots)}{\partial v_i}^2 \sigma_{v_i}^2}$$

The expression of f is

$$f \doteq (A^T A)^{-1} A^T B$$

where A is a matrix $n \times 3$ and B is a matrix of $n \times 1$. They are constructed with linear equations. These matrices depend on (u_i, v_i). We can infer a good criterion for checking if an edge is an occluding edge or not. Since a normal edge is characterized by a zero radius, the criterion is based on the probability for zero to be in the interval of confidence: $r - 2\sigma < 0 < r$.

3.2 Results of the implemented system

We have tested the algorithm on synthetic and real data.

3.2.1 Synthetic data

The test on synthetic data aims at testing the software and the verification of the noise model that we have used.

The principle of this test is to take a description of a system of real cameras and to simulate the observation of a cylinder. In fact, we only compute the image-segment of the extremal boundary of the cylinder. We add some noise to the endpoints of this image-segment. We use a Gaussian noise with a variance of one pixel.

The next table shows the value of the following parameters for a set of five cylinders: the radius of the cylinder, the measured radius, the value of the uncertainty σ and the criterion $\mathcal{C}$.

Real Radius	Estimated Radius	σ	$\mathcal{C}$
0	1.9	24.5	< 0
50	59.2	23.2	0.21
100	98.3	22.1	0.55
130	114.5	21.3	0.62
160	175.2	20.0	0.77

The baseline is approximately 250 millimeters wide and the distance from the optical center of the cameras to the objects is about 800 millimeters.

Nonetheless, we have to keep in mind that the uncertainty on the calibration of the cameras and the determination of the motion have not modelized.

3.2.2 Real data

For the test on real data, we have used small toys. The results are presented in figure (2). On the first image, we show a tea box, a mug and a cylinder with on the left side a tape. We then show two images:

- On the first one, the width of the line depends on the value of the criterion $\mathcal{C}$.

- On the second one, the width of the line depends on the value of the curvature radius.

The dotted lines correspond to segments which have not been matched between at least three images.

We can consider that we have detected an extremal boundary if the radius is non-null with a criterion which is sufficient. The results are satisfactory for the longer segments. There are occasional mistakes for the smaller segments.

One can note for example, that the segment-image of the cylinder are characterized by an important radius and a great value for the criterion $\mathcal{C}$.

The vertical segments are in general more precise because they have a better orientation with respect to the epipolar lines.

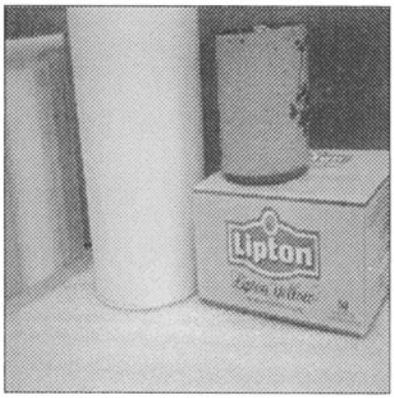

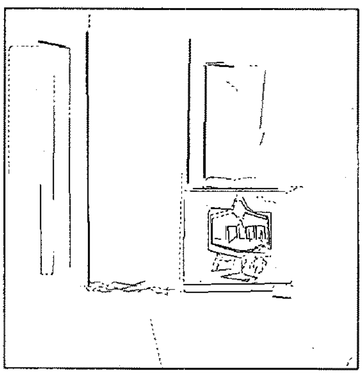

Scene 1, mode 1

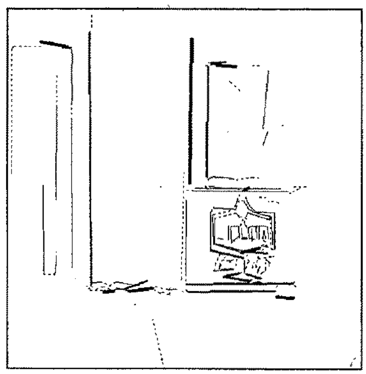

Scene 1, mode 2

Figure 2: Scene 1

3.3 Conclusion

The main output of this first part of our work is that it is possible to detect the extremal boundaries. We have to use a model of uncertainty. It is clear that this process can provide false results if the observed occluding contour does not correspond to a cylinder.

The following algorithm handles this case by dealing locally with the occluding contour results are corrected by the algorithm which deals locally with the occluding contour. Supplementary details about the cylinder case can be found in [8].

4 Estimating the object curvature along an occluding edge

In this part, we suppose that we have detected an extremal boundary and we want to compute some properties of the surface in the neighbourhood of the rim. We are interested by the differential properties of order up to two of the surface. Fundamental theorems of Differential Geometry [2] assert that these properties are sufficient to characterize the surface.

- The zero order differential property is the simple estimation of the position of the point. It means that we have to compute the exact position of the contact point between the surface and the optical ray for each of the cameras.
- The first order differential property is the estimation of the tangent plane to the surface. It is the easiest to obtain as we are observing an extremal boundary. In this case the tangent plane is the optical plane.
- The second order properties are the more difficult to obtain as they require the evaluation of second order derivatives. Such computation can be sensitive to noise.

4.1 Estimation of the position of the points

We first notice that the rim (R) of a surface is a curve, and thus the image of the rim (r) must be a curve. It is always true in a generic position. So, we can suppose that we have detected a curve (r_i) in each image. For each of these curves, it is possible to compute the tangent vectors at each of their points. The key idea is to neglect the apparent curvature and to use only the radial curvature. This can be realized by estimating the osculating circle of the radial curve. We have explained in [7] how this estimation can be carried out.

4.2 Computation of the second differential properties of the object shape

In this part, we show that it is possible to compute the two main curvatures of the surface near the points M_i.

4.2.1 Some useful equations

We have previously established that if we observe an occluding boundary and that if we suppose that (θ, ϕ) is an admissible parametrization of the surface in the neighborhood of M_i, located on the rim, we can obtain $X = X(\theta, \phi)$ from equations (1 - 3).

The solution of these equations gives us the analytical expression of $X(\theta,\phi)$. The two fundamental quadratic forms of a surface, which is represented by an admissible parametrization, can be derived from this expression.

At the end of all these computations we obtain that the evaluation of the first and second fundamental quadratic forms requires an estimation of the value of θ, ϕ, $p(\theta,\phi)$, $\frac{\partial p(\theta,\phi)}{\partial\theta}$, $\frac{\partial p(\theta,\phi)}{\partial\phi}$, $\frac{\partial^2 p(\theta,\phi)}{\partial\theta^2}$, $\frac{\partial^2 p(\theta,\phi)}{\partial\phi^2}$, $\frac{\partial^2 p(\theta,\phi)}{\partial\theta\partial\phi}$.

This is very interesting since these values can be estimated with sufficiently good accuracy for the points belonging to an extremal boundary. In [7], we detail the estimation of these values.

4.2.2 A particular case: zero Gaussian curvature

The previous algorithm is exact under the assumption that (θ,ϕ) is an acceptable parametrization of the observed surface in the neighborhood of M_i. For a generic position, this assumption fails if and only if one of the two main curvatures is equal to zero or in other words if the Gaussian curvature is equal to zero. But this is precisely the case where our cylinder model yields directly the answer: the computed radius of the cylinder gives us the first main curvature and the second is zero. The detection of this situation is performed directly by testing the curvatures of the image curves.

5 Experimental Results

We have tested the algorithm mostly on synthetic data. In fact we should say "almost" synthetic since, even though we have been using synthetic models (and one real ball), their rims have been projected on real 512 by 512 images and quantization noise is therefore present in the data. The reason why we have used synthetic images at this stage is that computing the curvatures requires calculating the second order derivatives and the process of differentiation is well known to be noise sensitive. Hence we have decided to isolate the possible sources of error by testing the algorithm on synthetic images first and to delay the experiments using real objects after we gain a better understanding of the numerical stability of the algorithm.

We have conducted a first set of experiments on synthetic images corresponding to a torus and a one-sheet hyperboloid. In order to analyze the results, we have focused on two things, the reconstructed points and the estimate of the curvatures.

Figures (3 - 5) represent the reconstruction obtained for a synthetic torus, a one-sheet hyperboloid and the image of a real volley-ball ballon. The figures show the points which are reconstructed from each camera, and the position of the centers of the circles found when applying the cylinder method.

We notice that we have reconstructed three different chains. On some parts, they are a bit noisy. These parts correspond to points where the epipolar plane (C_i, m_i, C_j) is tangent to the rim (r_j). In this case, the cylinder model fails to reconstruct the rim (R_i). This is not a problem of the method, but a general problem: there are two images which provide the same information. This situation is detected by the algorithm through the test on the variance on the radius of the cylinder.

The other figures represent the computed radii of curvature for the points which belong to the reconstructed rim. These curves are a little more difficult to interpret. Two of them

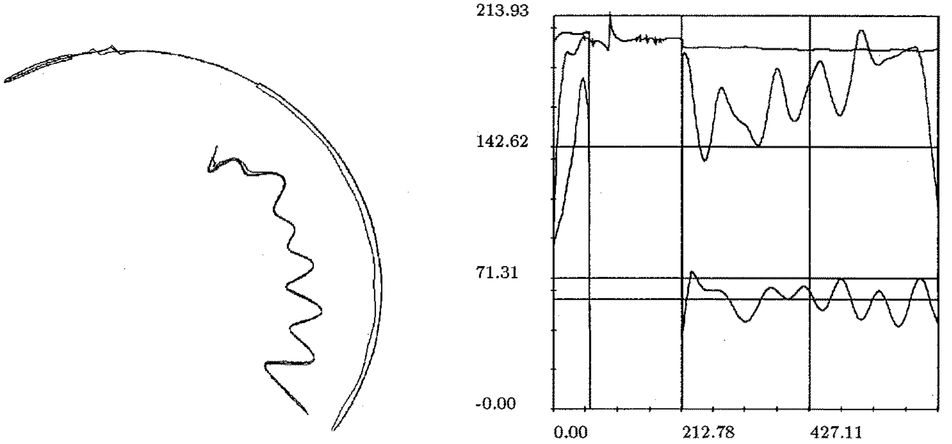

Figure 3: A synthetic torus

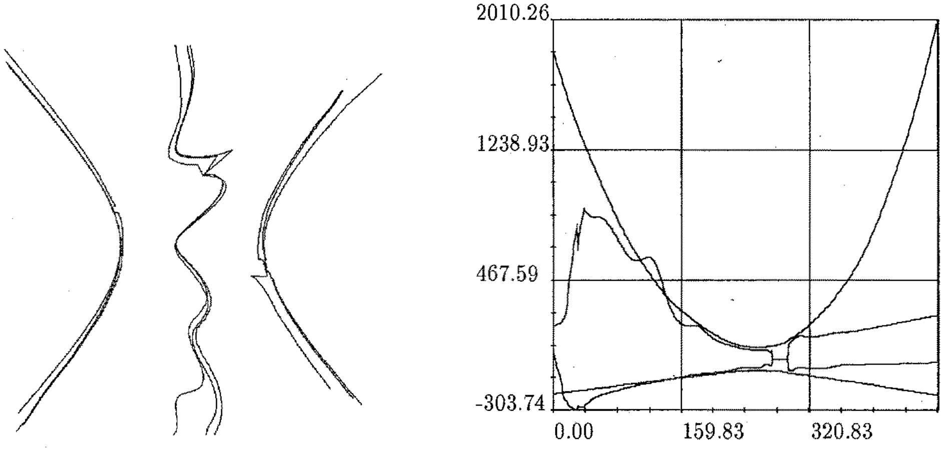

Figure 4: A synthetic one sheet hyperboloid

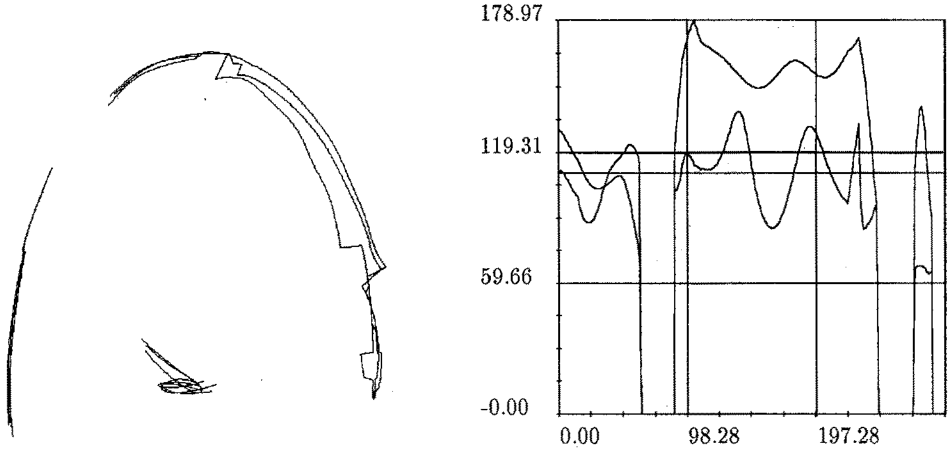

Figure 5: A real volley-ball ballon

correspond to the values of each of the two main curvatures along the extremal boundary. The two other curves correspond to the "theoretical" radii of curvature. Several parts can be distinguished:

- The two curvatures are equal to zero. For these points, the algorithm fails because of the epipolar plane problem.
- One of the two curvatures is equal to zero. This means that we have detect that the Gaussian curvature was equal to zero. This detection is obtained by testing the curvature of the image curve. In this case, the other curve in the figure shows the radius of the cylinder.
- The two curvatures are different from zero. They can be of the same sign (a positive Gaussian curvature) as in the case of the volley-ball or the synthetic sphere. The corresponding points on the surface are elliptic. They can have different signs (a negative Gaussian curvature) as in the case of the synthetic one-sheet hyperboloid. The corresponding points are hyperbolic.

6 Conclusion

In this article, we have shown that occluding edges were a robust source of 3D information. Points on the rim can be accurately reconstructed and good estimates for the second order differential properties of the surface in the vicinity of the rim can be reliably computed. More work needs to be done to test our algorithms further on a larger variety of shapes, study degenerate cases, and include this kind of processing in the framework of an active exploration of an object shape. We are actually testing our algorithm on a large number of images representing several different shapes of occluding edges. We also want use the algorithm with more than three views. The supplementary views will be obtained by moving the object with a known motion. We think that this will improve accuracy greatly.

Acknowledgements

We would like to thank Pr Olivier D. Faugeras and Michel Schmitt for their many helpful conversations and suggestions and Luc Robert for drawing the figures of this article.

References

[1] M. Brady, J. Ponce, A. Yulle, and H. Asada. Describing Surfaces. In Cambrdige MA USA MIT Press, editor, *Proceedings 2nd International Symposium Robotics Research*, 1986.

[2] M. P. Do Carmo. *Differential Geometry of Curves and Surfaces.* Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1976.

[3] Peter Giblin and R. Weiss. Reconstruction of Surfaces from Profiles. In *First International Conference on Computer Vision*, december 1986.

[4] D. Hilbert and S. Cohen-Vossen. *Geometrý and the Imagination.* Berlin : Springer, 1932.

[5] Jan Koenderink. What Does the Occluding Contour Tell Us About Solid Shape? *Perception*, 13:321–330, 1984.

[6] Basri Ronen and Ullman Shimon. The Alignement of Objects with Smooth Surfaces. In *Second International Conference on Computer Vision*, december 1988.

[7] Régis Vaillant and Olivier Faugeras. Using Extremal Boundaries for Recrovering Shape Properties of Objects. In *IEEE Workshop on Interpretation of 3D Scenes*, pages 26–32, november 1989.

[8] Régis Vaillant and Olivier Faugeras. Using Occluding Contours for 3D Object Modelling. In *International Advanced Robotics Programme*, october 1989.