Accurate Surface Reconstruction from Apparent Contours

Federico Pedersini, Augusto Sarti, Stefano Tubaro
Dip. di Elettronica e Informazione (DEI), Politecnico di Milano
Piazza L. Da Vinci 32, 20133 Milano, Italy
E-mail: pedersin/sarti/tubaro@elet.polimi.i

Abstract

In this paper we propose a 3D reconstruction technique based on the analysis of the extremal boundaries of the imaged surface. The method correctly determines the correspondences between apparent (viewer-dependent) contours and accurately determines location, orientation and curvature of the surface in the proximity of the apparent rims.

1 Introduction

The modeling of complex 3D scenes from the analysis of luminance edges using multi-camera systems can be a task of great complexity when some contours are generated by smooth self-occlusion of object boundaries. These extremal boundaries, in fact, fail to generate stereo-corresponding edges as their 3D coordinates are viewer-dependent. On the other hand, apparent contours [1] are known to provide valuable information on the local 3D structure of the surface (position, tangent plane and curvature) near the visible rims of the objects, provided that at least three views be available.

Several methods have been developed for estimating the 3D structure of the surface near the extremal boundaries [1, 2]. One solution [2] consists of locally approximating the surface with an osculating quadric, which is always defined except when the motion of the camera occurs along the viewing direction. Another solution [1] consists of reconstructing the "radial sections" of the object. The analysis, in this case, is performed on a family of planes called "radial planes", which are defined, for each point of the contour, by the visual ray and the normal to the contour on the image plane. In order to do so, three lines that are known to be tangent to the apparent rim are determined on the radial plane and the circle that is tangent to them (*osculating circle*) is taken as an local approximation of the radial section.

The choice of sectioning the object's surface with radial planes [1, 3] for 3D reconstruction purposes is not the only one possible. In fact, in this paper we propose a solution that uses the epipolar planes associated to one of the three possible pairs of views of the trinocular acquisition. This choice has the immediate advantage of simplifying the determination of the three lines that are tangent to the surface at the apparent rims, as two of them are already available (visual rays). We propose a method for determining the intersection between the rim associated to the third view and the epipolar plane through an exact geometrical procedure. Once the three points of tangency of the osculating circle are available on the epipolar plane, we can accurately model the radial section. As this procedure can be repeated for all the epipolar sections of the surface about the extremal rims, the final result is a complete parametrization of the self-occluding portion of the surface about and beyond the plane of visibility.

One other important reason for using epipolar sections instead of radial sections, is that it allows us to solve the problem of the determination of the correspondences between apparent contours, which has always been left aside in the literature. In fact, the available reconstruction methods based on extremal boundaries rely on the assumption that the contour correspondences are already available and correctly determined. As the stereometric principles cannot be applied to apparent contours, it would seem that we cannot rely on strong constraints for their correct matching. As a matter of fact, this is not entirely true, as there are points on the extremal boundaries where stereometric principles hold true. Such points correspond to the intersection between extremal contours on the surface to be reconstructed. This fact was used in the literature for generalizing the concept of epipolar constraint [5].

In this paper we exploit the generalized epipolar constraint and use it to validate correspondences between apparent contours. We also show that, when the extremal contours do not intersect on the viewed surface, a number of rules of congruence, consistency and smoothness, can be successfully used to correctly match apparent contours.

2 Contour matching

As homologous apparent contours do not back-project onto the same 3D object feature, we need to re-define the concept of correspondence. We will consider two apparent contours as corresponding if they smoothly merge into one another as the viewpoints tend to coincide.

With reference to Fig. 1, let us consider the two apparent contours $\mathbf{p}^{(1)}(s)$ and $\mathbf{p}^{(2)}(s)$ that are visible from viewpoints 1 and 2, respectively. Although such contours correspond to two different rims, r_1 and r_2, on the object's surface, their intersection $\mathbf{P}_{12}$ is still a point of stereo-correspondence. In fact, there exists an epipolar plane of the two views that is tangent to the surface at $\mathbf{P}_{12}$. This epipolar tangency constraint can be formulated as follows: the line that passes through the epipole $\mathbf{e}_2^{(1)}$ and is tangent to the apparent contour on image 1 is projectively related to the line that passes through the epipole $\mathbf{e}_1^{(2)}$ and is tangent to the apparent contour on image 2 [5].

When considering three views, the extremal rims, r_1, r_2 and r_3, on the object's surface intersect each other only pairwise. However, we notice that the three portions of the rims between the points of intersections (from $\mathbf{P}_{31}$ to $\mathbf{P}_{12}$ on rim r_1, from $\mathbf{P}_{12}$ to $\mathbf{P}_{23}$ on rim r_2, and from rim $\mathbf{P}_{23}$ to $\mathbf{P}_{31}$ on rim r_3) form a closed piecewise-smooth curve. As $\mathbf{P}_{12}$, $\mathbf{P}_{23}$, and $\mathbf{P}_{31}$ are points of stereo-correspondence, we can "follow" such a closed curve on the three views by tracking each one of the apparent contours, and "jumping" from view to view whenever we meet an epipolar tangency point $\mathbf{p}_i^{(j)}$. The correspondence between three apparent contours is guaranteed when we can follow a cyclic "closed" path along the apparent contours using the epipolar tangency points as points of connection between views.

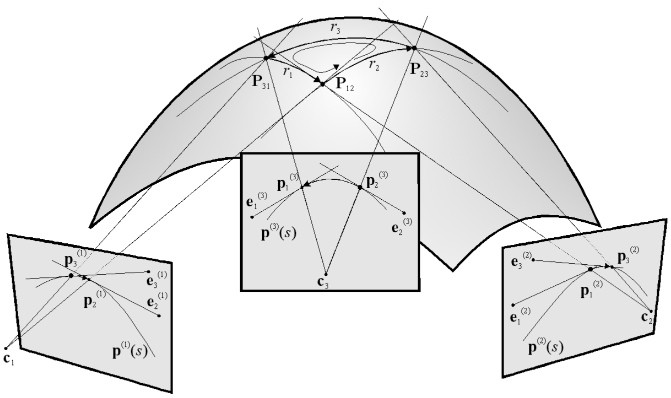

Figure 1: Illustration of the generalized trinocular constraint.

In order to make sure that the matching algorithm will perform correctly even when it is not possible to determine two epipolar lines that are tangent to each one of the candidate homologous contours, we adopt an optimization strategy based on an appropriate objective function, which we apply to a reduced search space of potential matches. A significant reduction of the search space is a crucial step of the matching strategy and can be achieved using smoothness constraints and rules of congruence and consistency. The global matching algorithm can thus be summarized as follows:

1. build a list of candidate contours;
2. rule out sharp contours;
3. find exact correspondences;
4. rule out inconsistent correspondences;
5. select remaining matches through optimization.

Step 2 is performed by removing stereo-corresponding edges from the list. Step 3 is carried out using the generalized trinocular constraint. Step 4 exploits the fact that extremal surfaces cannot cover visible details, intersect sharp 3D contours or other extremal surfaces. Furthermore it rules out correspondences that fail to produce smooth apparent rims and/or extremal surfaces. Finally, Step 5 is based on the minimization of an appropriate cost function that we measures the "inconsistency" between the orientations of the apparent rims. In order to do so, we compute the angles between the viewing planes that are tangent to the rims and the epipolar plane that results as being most orthogonal to the apparent contours. In fact, it is quite reasonable to expect that such angles do not significantly differ from one another at the intersections between rims and epipolar plane. Such intersections can be determined as explained in the next Section.

3 Reconstruction on epipolar sections

As already said above, the 3D reconstruction is performed on epipolar planes, which are always selected in such a way to be as orthogonal as possible to the apparent contour. Given a triplet of matched contours, we scan each one of them with the most proper epipolar plane and, for each epipolar section, we determine the three points of intersection with the apparent rims (see Fig. 2). In order to do so, we first determine the three lines on the epipolar plane that pass through such points and are tangent to the epipolar section of the surface. The points of intersection with the three rims will be those of such lines that belong to their osculating circle.

Let us consider a matched triplet of apparent contours (c_1, c_2, c_3). Given a point of c_1, there are two possible epipolar planes, depending on which other camera we are considering. Let π_e be the one that is orthogonal the most to c_1. The goal is to determine the three points p_1, p_2 and p_3 of intersection between π_e and the apparent rims r_1, r_2 and r_3 on the object surface $\mathcal{S}$. In order to do so, we first determine the three lines v_1, v_2 and v_3 on π_e that pass through p_1, p_2 and p_3, respectively, and are tangent to epipolar section $s_e = \mathcal{S} \cap \pi_e$ of the object's surface. Such points of intersection can be used to determine the osculating circle that approximates the epipolar section of the surface.

As we can see in Fig. 2, two of such tangents are are already available as they are the two visual rays v_1 and v_2 that generated the epipolar plane π_e. We thus need to determine the third tangent r_3. In order to do so, let us consider the plane π_t that is tangent to the third rim at one of its points q and that passes through the optical center c_3 of the third camera. In order to determine the third tangent v_3, we will analyze the behavior of the line v_t of intersection between π_t and π_e. In fact, as q slides along the third rim, the radius of the osculating circle s_t to v_1, v_2 and v_t changes, and exhibits an extremum when $v_t = v_3$, i.e. when $q \in \pi_e$ (which means that $q = p_3$). Whether this extremum is a minimum or a maximum depends on the geometry of the intersection of the tangents and on the sign of the curvature of the apparent contour c_1.

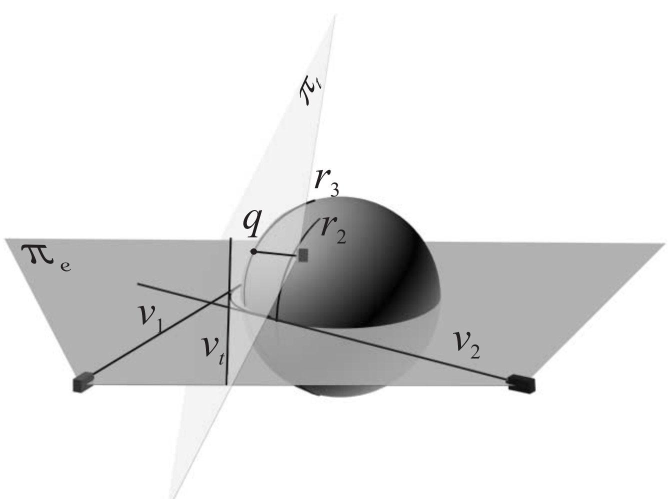

Figure 2: Reconstruction of the epipolar section. Two tangent lines are already available (lines of sight). The third one is obtained as the intersection between the epipolar plane and the tangent plane to the third rim that gives rise to an osculating circle of maximum or minimum radius, depending on the geometric configuration of the three lines and the sign of the curvature of the horizon contour.

For example, let us assume that the reference horizon contour c_1 is convex, and let us consider a circular epipolar section s_e of radius R, and the geometry of intersection of its tangents depicted in Fig. 3*a*. If v_t falls outside of s_e, then the osculating circle to v_1, v_2 and v_t has a radius R_t that is smaller than R. More precisely, if v_t intersects v_1 between A and B, and v_2 between A and C, then $R_t < R$. If the geometry of the intersection between v_1, v_2 and v_t is different from that of Fig. 3*a* (see, for example, Fig. 3*b*), then the osculating circle to v_1, v_2 and v_t has a radius R_t that is greater than R. Similar geometrical considerations apply to the case of convex rims, but with opposite signs of the inequalities.

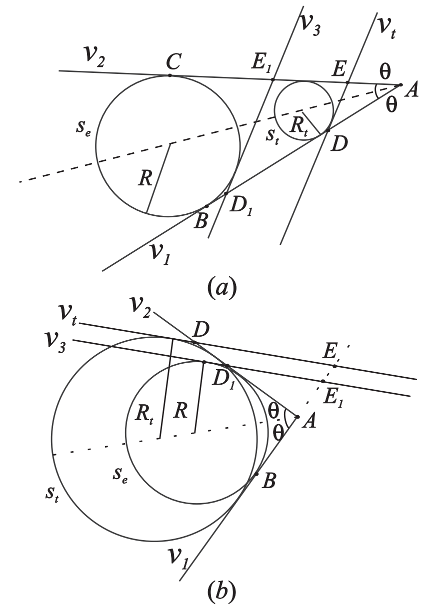

Figure 3: As the tangent plane π_t slides along a convex rim, the line $v_t = \pi_t \cap \pi_e$ tends either to drift away (a) from the epipolar section s_e, or to cut through it more and more (b), depending on the geometry of the intersection between tangents.

In conclusion, in order to determine v_3, we will analyze the behavior of the radius R_t of the osculating circle to v_1, v_2 and $v_t = \pi_t \cap \pi_e$, as the tangent plane slides along the third rim in the proximity of π_e. The search for a maximum or a minimum of R_t, is performed in closed form according to the above geometrical considerations.

It is important to notice that, unlike the methods proposed in [1, 3], this method determines the three tangent lines through an exact geometrical procedure. Furthermore, there are no restrictive hypotheses on the local shape of the surface. The only assumption that we make is that the curvature of the apparent contour does not change sign, which is not restrictive if we accordingly pre-segment the contours.

The above approach is used for determining the three points p_1, p_2 and p_3 of intersection between π_e and the horizon rims. And this is done for all points of the reference horizon contour, and for all three contours of the matched triplet. Once the triplets of points of intersection have been computed for all epipolar sections, we can proceed with the final 3D reconstruction as we prefer. For example, we can use just the 3D coordinates of such points, or use the fact that the object surface is bound to be tangent to the visual rays or, finally, derive the local surface curvature from that of the osculating circle for all epipolar sections.

4 Experimental results

An example of 3D reconstruction on epipolar sections of an object containing only apparent contours (a torus) is shown in Fig. 4, together with an evaluation of the reconstruction's accuracy compared with that of a method based on radial sections [3]. Other examples of 3D reconstruction in the case of more complex scenes (containing both sharp and horizon contours) are shown in Fig. 5.

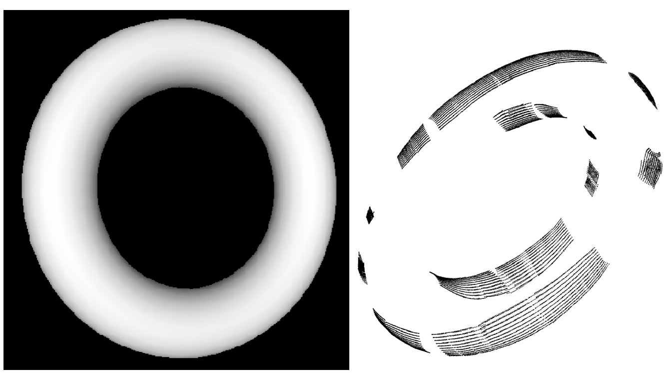

Figure 4: Example of reconstruction in the absence of sharp edges. The radii of the torus are r and $4r$. The average and the maximum reconstruction errors are $0.0118r$ and $0.0384r$, respectively, with a std. dev. of $0.0089r$. The 3D reconstruction on radial planes results in average and maximum reconstruction errors of $0.0130r$ and $0.0988r$, respectively, with a std. dev. of $0.0110r$.

5 Conclusions

The method has proven to be capable to correctly matching, classifying and reconstructing 3D contours in a completely automatic fashion, with considerable accuracy and with virtually no outliers.

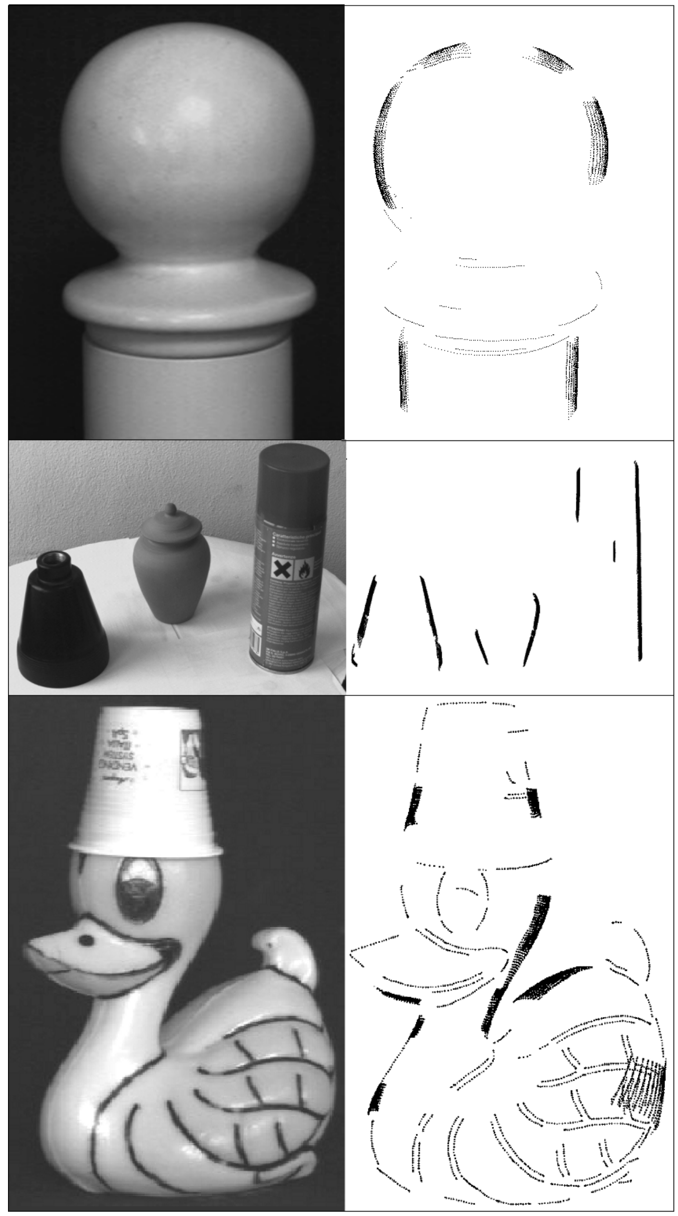

Figure 5: Other examples of reconstruction with sharp and horizon edges. Contour matching is completely automatic.

References

[1] R. Vaillant, O.D. Faugeras, "Using extremal boundaries for 3D object modeling". *IEEE Tr. PAMI*, Vol. 14, No. 2, Feb. 1986, pp. 157-173.

[2] E. Boyer, M.O. Berger, "3D surface reconstruction using occluding contours", *Intl. J. of Computer Vision*, Vol. 22, No. 3, pp. 219-233, 1997.

[3] F. Pedersini, A. Sarti, S. Tubaro, "3D surface reconstruction from horizons". IWSNHC3DI, Rhodes, Greece, 1997.

[4] K. Åström, R. Cipolla, P.J. Giblin, "Generalised epipolar constraints", *Proc. ECCV*, Vol. 2, p. 97-108, Cambridge, UK, 1996.