Reconstruction of Canal Surfaces from Single Images Under Exact Perspective

Vincenzo Caglioti and Alessandro Giusti

Politecnico di Milano
{caglioti, giusti}@elet.polimi.it

Abstract. This paper addresses the reconstruction of canal surfaces from single images. A canal surface is obtained as the envelope of a family of spheres of constant radius, whose center is swept along a space curve, called axis. Previous studies either used approximate relationships (quasi-invariants), or they addressed the recognition based on a geometric model. In this paper we show that, under broad conditions, canal surfaces can be reconstructed from single images under exact perspective. In particular, canal surfaces with planar axis can even be reconstructed from a single fully-uncalibrated image. An automatic reconstruction method has been implemented. Simulations and experimental results on real images are also presented.

1 Introduction

One of the prominent problems in computer vision is reconstruction of the shape of 3D objects from a single, bidimensional image. This work, in particular, deals with a *shape from contour* problem: reconstruction of a canal surface from a single perspective image. A canal surface is obtained as the envelope of a family of spheres of constant radius, whose center is swept along a space curve, called axis.

Circular cross section pipes and flexible wires can be modeled as canal surfaces, and reconstructed with this approach. Moreover, long-exposure photographs of a moving sphere (e.g. a kicked soccer ball) are images of canal surfaces as well, therefore we are also applying this technique to sport environments in order to analyze particular nonparabolic trajectories deriving from fast ball spin.

Some approaches about shape reconstruction of such objects are based on information about the surface normal [1], other approaches consist of shape from shading techniques based on Lambertian model [2]. Some other approaches are based on the use of stereoscopic vision [3]. Approximate relationships, as, e.g., quasi-invariants, are used in [4, 5, 6]. A reconstruction method for orthogonal projections, which requires that at least a cross section is visible, is presented in [7].

Other publications ([8, 9]) focus on the geometric properties of generalized cylinders, but do not deal with the reconstruction process. Likewise, works such as [10] aim at identifying the 2D perspective projection of the axis of revolution

A. Leonardis, H. Bischof, and A. Pinz (Eds.): ECCV 2006, Part I, LNCS 3951, pp. 289–300, 2006.

and do not return a full 3D reconstruction of the shape. An example of full 3D reconstruction of another class of generalized cylinders has been presented in [11], that deals with solids of revolution.

This work is entirely based on the geometric properties of canal surfaces and of their apparent contours in perspective images, and allows to find a full 3D reconstruction of the canal surface and its curvilinear axis. In particular, canal surfaces with planar axis can even be reconstructed from a single fully-uncalibrated image, while nonplanar-axis canal surfaces need a calibrated image. [12] provides a useful algorithm for contour tangent direction estimation.

Section 2 provides some basic definitions and properties, which are used in section 3 to derive the key relations used in the paper; section 4 describes how we deal with uncalibrated images; section 5 details the geometric considerations driving the actual reconstruction process; section 6 conveys a broad view of the complete reconstruction process, whereas section 7 describes the results obtained by our prototype implementation. Section 8 presents the conclusions and future directions of our work.

2 Definitions and Basic Properties

A canal surface can be defined as the envelope surface of a family of spheres with constant radius R, whose centers lie on a space curve called axis, such that, at any axis point, the axis curvature radius is strictly larger than R.

A planar-axis canal surface is a canal surface whose axis is a planar curve.

Property 1. A canal surface is equivalent to the union of circumferences with radius R, called cross sections, such that each cross section is centered on the axis. An axis point and the cross section centered on it are said to be associated. A cross section has a supporting plane perpendicular to the tangent to the axis at its associated point.

A canal surface projects a pair of facing apparent contours; our approach only considers the lateral contours of the canal surface, and does not require any cross section to be visible.

Two contour points are said to be coupled if they are the image of two points on the same cross section.

Property 2. Let P be a point on the canal surface, C be the cross section on which P lies, and P_s the associated axis point: Let T be the tangent plane to the canal surface in P: T is parallel to the tangent to the axis at P_s.

An immediate consequence is that the tangent plane is perpendicular to the plane supporting C.

For any point on the axis P_s, we can define a Tangent Cylinder (TC): The TC has radius R and axis tangent to the canal surface axis in P_s. The intersection between the TC and the canal surface contains the cross section centered on P_s. If the axis is rectilinear, the TC coincides with the canal surface.

3 Properties of Canal Surface Contours

A number of properties of the apparent contours are presented in this section.

3.1 Coupling Condition

First, we present a necessary condition[1] for the coupling of contour points which involves the camera parameters, but holds regardless of the geometry of the axis. The property is used to detect coupled points on contours, enabling us to reconstruct the axis shape when the camera parameters are known; it is also used in the opposite direction, generating constraints for camera parameters when two coupled points are known in advance: this allows to calibrate the camera when contour features presented in the following allow to detect pairs of coupled points.

Property 3. Let c_1 and c_2 be two facing contours on the image; let t_1 (t_2) be the tangent to c_1 (c_2) at point p_1 (p_2), and let v_h be the intersection between t_1 and t_2.

The points p_1 and p_2 are coupled only if the angle formed by $\overline{Op_1}$ and $\overline{Ov_h}$ coincides with the angle formed by $\overline{Op_2}$ and $\overline{Ov_h}$, where O is the camera viewpoint.

Proof. Now we prove the above necessary condition.

Let P_1 (P_2) be the point on the canal surface which projects to p_1 (p_2), let T_1 (T_2) be the plane tangent to the canal surface at P_1 (P_2); note that T_1 (T_2) is the interpretation plane of t_1 (t_2).

Let C be the cross section containing P_1 and P_2, and let P_s be the axis point, center of C. Let Π_{sym} be the plane bisecting T_1 and T_2: since both T_1 and T_2 are tangent to C (which is a circumference) and perpendicular to its supporting plane, C is symmetrical w.r.t. Π_{sym}; P_1 and P_2, intersection of symmetrical entities, are symmetrical as well. Let V be the intersection line between T_1 and T_2: V lies on Π_{sym}; The camera viewpoint O, which belongs to V, lies on Π_{sym} as well. Therefore, the angle formed by $\overline{OP_1}$ and V equals the angle formed by symmetrical entities $\overline{OP_2}$ and V. The thesis immediately follows.

This condition is necessary but not sufficient for the coupling of p_1 and p_2: however, if p_2 is constrained to lie on a curve c_2, and t_2 is constrained to be tangent to c_2 in p_2, few, sparse choices of p_2 satisfy the condition[2].

3.2 Properties of Planar-Axis Canal Surfaces

When the axis of the canal surface is constrained to lie on a plane, additional properties hold.

[1] See [10] for a similar property for surfaces of revolution; note that its extension to canal surfaces is not straightforward.

[2] An exception is the degenerate case where parts of c_2 coincide to an arc of an ellipse which is both tangent to t_1 in p_1 and image of a sphere.

In the following, three possible relations between the viewpoint position and the canal surface will be considered: consider the two planes parallel to the axis plane, at distance R from it; they are tangent to the canal surface at two diametric points for each cross section, and the canal surface is entirely enclosed between the two planes; the viewpoint can lie outside the space enclosed by the two planes (*configuration 1*), between the two planes (*configuration 2*), or on one of the two planes (degenerate *configuration 3*). If the axis is rectilinear, the canal surface is a cylinder, therefore the following considerations do not apply.

Inflection Points

Property 4. If a planar-axis canal surface is seen by a camera whose viewpoint is placed according to configuration 1 or configuration 2, an inflection point on one contour is always coupled to an inflection point on the facing contour, and the related axis point is an inflection point for the axis.

A proof is given in [6]; note that this property is independent of camera calibration.

3.3 Bitangents

Inflection points on contours are not the only useful feature: also bitangents to canal surface contours allow to determine coupled points regardless of the camera calibration parameters, by means of the following property (see figure 1):

Property 5. Let b_1 be a bitangent to contour c_1, and name p_1^a and p_1^b the two tangency points. If the viewpoint is placed according to configuration 1 or configuration 2, a bitangent (b_2) to the contour c_2 exists, and its tangency points p_2^a and p_2^b are coupled with p_1^a and p_1^b respectively[3].

Proof. Let P_1^a, P_1^b, P_2^a and P_2^b be the points on the canal surface which project to p_1^a, p_1^b, p_2^a and p_2^b respectively; let C^a be the cross section passing through by P_1^a and P_2^a, and C^b the cross section passing through P_1^b and P_2^b; call P_s^a (P_s^b) the axis point at the center of C^a (C^b), and D^a (D^b) the directions of the tangent to the axis in P_s^a (P_s^b).

Let T_1 be the interpretation plane of b_1: T_1 is tangent to the canal surface in P_1^a and P_1^b; then, T_1 contains both D^a and D^b. Moreover, D^a and D^b are constrained to be parallel to the axis plane. Since T_1 and the axis plane are not parallel, D^a and D^b coincide; therefore C^a and C^b lie on parallel planes.

Because C^a and C^b are two circumferences tangent to the same plane (T_1) and lying on parallel planes, their centers P_s^a and P_s^b lie on a plane parallel to T_1; moreover, being axis points, P_s^a and P_s^b must lie on the axis plane. Let Λ be the line connecting P_s^a and P_s^b: since T_1 and the axis plane are not parallel, Λ has direction $D^a = D^b$; Λ is a bitangent for the axis, with P_s^a and P_s^b as

[3] The property requires minor adjustments to deal with spines whose tangent orientation varies broadly.

tangency points; moreover, since C^a and C^b lie on planes perpendicular to Λ, they are cross sections of the right cylinder T_{cyl}, which is the tangent cylinder to the canal surface in both C^a and C^b.

Let T_2 be the other plane, besides T_1, tangent to T_{cyl} and passing through the viewpoint O: T_2 is also tangent to the canal surface in P_2^a (which belongs to C^a) and P_2^b (which belongs to C^b); therefore, T_2 projects to a single line b_2, which is a bitangent for c_2 in p_2^a and p_2^b; moreover, p_1^a is coupled with p_2^a, because they are images of points belonging to the same cross section; similarly, p_1^b is coupled with p_2^b.

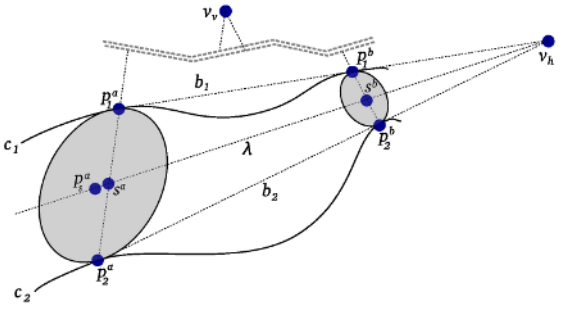

Fig. 1. Coupled bitangents and related vanishing points

Coupled bitangents also allow to find another constraint on camera calibration parameters: with relation to the entities defined above, the following property holds:

Property 6. Let v^a be the image line passing through p_1^a and p_2^a, v^b the image line passing through p_1^b and p_2^b, v_v the intersection of v^a and v^b, and O the viewpoint; let λ be the image of Λ; the direction identified by vanishing point v_v is orthogonal to the vector connecting O to any point on λ.

The property follows from the symmetry of P_1^a, P_1^b, P_2^a and P_2^b w.r.t. the plane containing Λ and O.

Three points on λ can be extracted from a a pair of coupled bitangents: v_h, intersection of b_1 and b_2; s^a, found using the cross ratio on p_1^a, s^a, p_2^a and v_v; and s^b, found similarly on the other cross section[4].

The maximum number of bitangents to a planar curve grows more than linearly with the number of inflection points: elaborate axis shapes are then likely to have a large number of bitangents; many meaningful bitangents can also be found bridging a number of canal surfaces with planar axis, which share the radius and axis plane: think of a set of identical torii placed on a planar surface.

Rectilinear parts on contours share the properties of bitangents; moreover, unlike bitangents they can also be exploited in the 3D-axis case.

[4] Note that s^a and s^b are not the images of the center of C^a and C^b.

4 Uncalibrated Camera

When an unknown canal surface is seen from an unknown camera, the contour properties presented in the previous section allow to define a number of constraints on the camera calibration parameters; if a sufficient number of constraints is defined, the camera can be calibrated, then the reconstruction can be carried out as detailed in the following section.

Although theoretically the camera could be calibrated even when the axis is not planar, provided that enough rectilinear parts on the canal surface contours allow to determine a sufficient number of coupled points pairs, we focus on the planar axis case. We can therefore use constraints originating from coupled inflection points on the contours and from coupled bitangents.

- Using property 3 in the reverse direction, a pair of coupled image points p_1 and p_2 enables us to enforce that

$$\frac{p_1^\top \omega v_h}{\sqrt{p_1^\top \omega p_1} \cdot \sqrt{v_h^\top \omega v_h}} = \frac{p_2^\top \omega v_h}{\sqrt{p_2^\top \omega p_2} \cdot \sqrt{v_h^\top \omega v_h}} \tag{1}$$

  where p_1 and p_2 are expressed in homogeneous coordinates, and ω is the image of the absolute conic, related with the calibration matrix K by $\omega = K^{-\top}K^{-1}$. The pair of coupled points can be identified on the image by means of property 4 or property 5.
- In addition to the equations presented before, according to property 6 a pair of coupled bitangents or rectilinear parts allows us to enforce the following linear constraints on ω:

$$s_a^\top \omega v_v = 0 \tag{2}$$

$$s_b^\top \omega v_v = 0 \tag{3}$$

$$v_h^\top \omega v_v = 0 \tag{4}$$

  where s_a and s_b have been defined in property 6. Two of these relations are independent.
- Regardless of inflection points and bitangents on contours, a valid camera calibration hypothesis allows a reconstruction where all found axis points and axis tangent directions lie on the same plane – the axis plane. The planarity of the axis tangent directions is easily checked by quantifying how well intersection points of coupled points' tangents fit to a line (the image of the line at the infinite of the axis plane). This constraint could be used in the absence of features such as bitangents or inflection points.

5 Canal Surface Reconstruction

For every pair of coupled points p_1 and p_2, the associated cross section in space can be reconstructed without ambiguity, provided that the radius of the canal surface is known.

The cross section orientation is represented by vanishing point v_h, intersection of tangents t_1 and t_2 (see figure 2). The angle α between T_1 and T_2 is computed; since the cross section radius R is known, and both T_1 and T_2 are tangent to C and perpendicular to its supporting plane, the distance between P_s and V is determined as a function of α alone. P_s is also constrained to lie on the plane bisecting T_1 and T_2. In conclusion, the cross section position is completely specified by constraining its tangency points to T_1 or T_2 to lie on the interpretation line of P_1 or P_2.

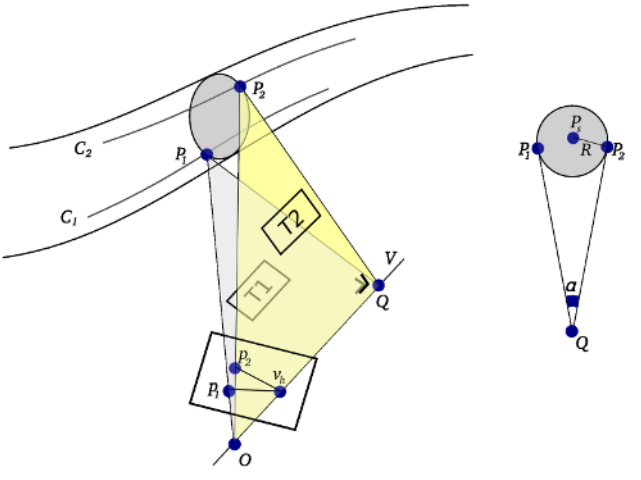

Fig. 2. Cross section reconstruction

Note that a scaled version of the canal surface can be reconstructed by using an arbitrary value for R.

When an ordered sequence of coupled point pairs along two facing contours is known, cross sections reconstructed from adjacent coupled point pairs can be joined in order to approximate the canal surface.

6 Implementation Notes

Starting from the input image, reconstruction requires to perform a number of sequential steps:

1. Edges on the input image are found using the Canny algorithm ([13]), and the edge points are localized with subpixel precision by fitting a gaussian curve to the gradient values around the found pixels; this allows to detect edges with enhanced precision.
2. Edge points are subdivided into chains, using an edge tracking algorithm biased towards smooth contours, which tolerates small discontinuities;

3. An estimate of the direction of the tangent to the contour in each of the edge points is computed, using the angle median method presented at [12];
4. If the camera calibration parameters are unknown and the axis is planar, inflection points and bitangents (section 6.1) are detected and coupled; the camera is then calibrated as described in section 4.
5. Contour points are coupled according to the procedures presented in section 6.3;
6. For each of the found couples, a cross section in space is reconstructed, exploiting the geometrical construction presented in section 5.
7. An optional postprocessing filter is used to mitigate the errors in the cross section localization.

6.1 Detecting Bitangents

Our input from the previous steps is a set of edge chains; each contour point is annotated with the orientation of the tangent to the contour, which is considered continuous along the contour[5] (except that around angular points).

A bitangent is defined between two edge points p^a and p^b if the contour tangent at p^a is collinear both with the contour direction at p^b, and with the direction of the vector connecting p^a to p^b.

Unfortunately, a threshold-based algorithm tends to detect clusters of many nearby bitangents if contours around tangency points have low curvature; therefore, we implemented an algorithm which filters out unwanted results, and has proved very effective in our tests:

1. A candidate bitangents list is populated with a threshold-based criterion;
2. The bitangents are ranked according to their alignment;
3. The highest ranked bitangent is extracted and returned, and all nearby bitangents are recursively discarded from the list; the step is repeated until the list is empty.

6.2 Coupling Inflection Points and Bitangents

Unpaired inflection points and bitangents are useless for camera calibration: we must determine which pairs of bitangents or inflection points are actually coupled, in order to determine the constraints presented at section 4.

Since the number of inflection points and bitangents in an image is usually limited, simple heuristics can be used in order to couple features; for example, a pair of coupled points must be near the tangency points of a circle, bitangent to the respective contours, confined inside the projection of the canal surface (i.e. not extending to overlap with the background)[6].

Moreover, several other rules based on simple geometric considerations allow to further constrain the possible solutions.

[5] Therefore, its range is not defined.

[6] Note that in the calibrated case we would be able to use an ellipse as the exact image of a sphere, instead of a circle.

6.3 Coupling Contour Points

As we noted previously, the condition stated in property 3 is a necessary condition for a pair of points to be coupled, but not a sufficient one: therefore, given a point on a contour, it does not usually allow to determine a single candidate coupled point, but suggests a set of possible candidates.

However, if a pair of coupled points is known, other pairs can be searched in the proximity on the facing contours: the two facing contours can also be given a consistent mutual orientation, in order to further reduce the search scope.

We define a fitness value $J_c(p_1, p_2)$ which quantifies how well a pair of contour points meets the condition stated in property 3, as the squared difference between the two angles $p_1\widehat{O}v_h$ and $p_2\widehat{O}v_h$:

The algorithm starts from an initial pair of coupled points, determined e.g. by property 3 or by a pair of bitangents, and, starting from this pair, other pairs are found incrementally by "walking" along the coupled contours, limiting the search of candidate coupled points to a very limited set at each iteration, and choosing the one minimizing J_c.

The processed contour parts are marked, then the algorithm is applied again with a different starting pair, ignoring contour parts already considered.

7 Experimental Results

We implemented the presented procedures in a Java-based prototype. Experimental results are presented with both simulated images and photographic ones; both planar-axis canal surfaces and nonplanar-axis canal surfaces are represented: several images with enough contour features for camera calibration have also been used.

To evaluate reconstruction results in photographic images, where the actual shape is not known, a preliminary qualitative evaluation has been carried out by

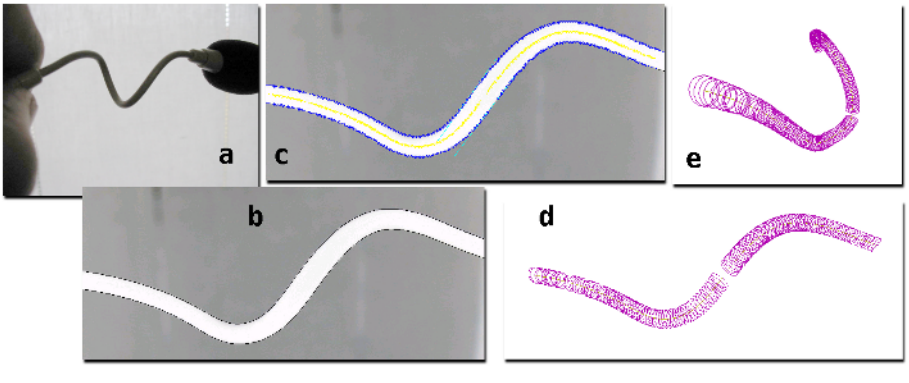

Fig. 3. Reconstruction of a 3D-axis canal surface from calibrated image: the object geometry (a), original image with edge detection (b), detection of coupled points and image of associated axis points (c), 3D view of reconstructed object (d, e)

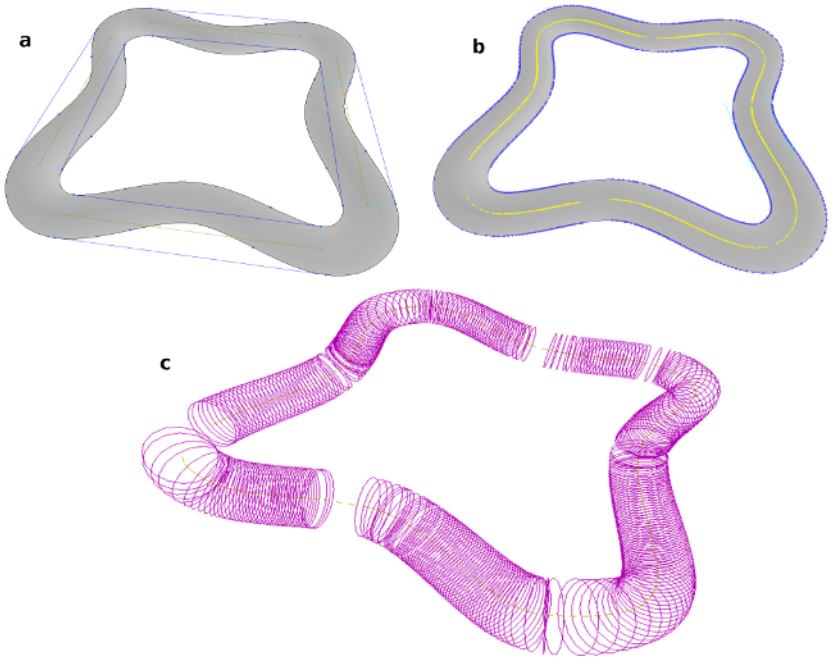

Fig. 4. A planar-axis canal surface (synthesized): original image with edge detection and detected bitangents and inflection points (a), detection of coupled points and image of associated axis points (b), 3D view of reconstructed object from different viewpoint (c)

reconstructing symmetrical canal surfaces seen from a generic viewpoint, then assessing the symmetry of the reconstructed shape; in this respect, we observed that the reconstruction is free of systematic errors; we also noted a remarkable robustness w.r.t. errors in given camera parameters.

Errors in placement and orientation of individual cross sections is heavily dependent on the quality of edges and on the distance between facing contours. We observed that when coupled points on facing contours are seen, from the viewpoint O, within an angle of more than 1/30 rad, errors in the localization of cross sections are acceptable; in particular, in the synthetic image in figure 4, where cross sections are viewed under an angle less than 1/10 rad, the average displacement error of the reconstructed cross section is within 1/10 of the cross section radius. Note that the effect of this error can be heavily mitigated by a moving average on neighboring cross sections, or more sophisticated techniques such as curve fitting. Also, the reconstruction quality heavily depends on the perspective effects: as these increase, the error decreases, and vice versa. It halves as the camera field of view is widened by 15°, it doubles as the canal surface radius is reduced by 1/3. In photographic images obtained with a standard 2Mpixel camera, we observed that the error variance increases by a factor between 2 and 4 w.r.t. a synthetic image with similar characteristics,

depending on contour sharpness, precision in camera calibration, and nonidealities of the imaged canal surface objects (unavoidable with flexible wires, for example).

Although not tuned for computational efficiency, the actual point coupling and surface reconstruction phases always required less time than the preceding edge detection and tracking steps. The whole procedure for the 3D-axis canal surface represented in figure 3 takes about 8 seconds on a Pentium 4 system, but only about 1.8 seconds are due to the actual point coupling procedure and reconstruction.

The planar-axis uncalibrated case has been tested as well.

The test image of a planar-axis canal surface depicted in Figure 4 allowed to calibrate the camera with an 8% average error by using only the linear constraints (2), (3), (4): since most bitangents' endpoints were affected by rather large localization errors along the contours due to minimal curvature around the tangency points, the determination of their v_v has been quite imprecise; the results improve by adding the nonlinear constraint (1), which is robust w.r.t. this sort of error.

8 Conclusions and Future Work

We presented a technique for reconstructing a canal surface from a single perspective image. The developed technique allows to reconstruct a canal surface, having a nonplanar axis, from a calibrated image. Moreover, canal surfaces with planar axis can be reconstructed from a single, fully-uncalibrated image. The implemented technique has been validated through experiments with both simulated and real images.

The present version of the full-uncalibrated reconstruction technique is based on projective-invariant features such as inflection points of bitangents: ongoing activity is aimed at the extension of this technique to cases, where such invariant features are not visible in the image.

References

1. J. R. Kender, R.K.: On seing spaghetti: self-adjusting piecewise toroidal recognition of flexible extruded objects. IEEE Transactions on Pattern Analysis and Machine Intelligence **17** (1995)
2. A. D. Gross, T.E.B.: Recovery of shcgs from single intensity views. IEEE Transactions on Pattern Analysis and Machine Intelligence **18** (1996)
3. J. M. Chung, T.N.: Extracting parametric descriptions of circular generalized cylinders from a pair of contours for 3-d shape recognition. In: Proc. of the 1995 IEEE Conference on Robotics and Automations. (1995)
4. M. Zerroug, R.N.: Three-dimensional descriptions based on the analysis of the invariant and quasi-invariant properties of some curved-axis generalized cylinders. IEEE Transactions on Pattern Analysis and Machine Intelligence **18** (1996)
5. M. Zerroug, R.N.: Part-based 3d descriptions of complex objects from a single view. IEEE Transactions on Pattern Analysis and Machine Intelligence **21** (1999)

6. N. Pillow, S. Utcke, A.Z.: Viewpoint-invariant representation of generalized cylinders using the symmetry set. In: Proc. of the British Machine Vision Conference. (1994)
7. F. Ulupinar, R.N.: Shape from contour. IEEE Transactions on Pattern Analysis and Machine Intelligence **17** (1995)
8. J. Ponce, D.C.: Finding the limbs and the cusps of generalized cylinders. International Journal of Computer Vision **1** (1987)
9. J. Ponce, D. Chelberg, W.B.M.: Invariant properties of straight homogeneous generalized cylinders and their contours. IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 11 **11** (1989)
10. R. Glachet, M. Dhome, J.T.L.: Finding the perspective projection of an axis of revolution. Pattern Recognition Letters **12** (1991)
11. Pernici, del Bimbo, C.: Metric 3d reconstruction and texture acquisition of surfaces of revolution from a single uncalibrated view. IEEE Transactions on Pattern Analysis and Machine Intelligence **27** (2005)
12. J. Matas, Z. Shao, J.K.: Estimation of curvature and tangent direction by median filtered differencing. In: Proc. of the 8 th Intl Conference on Image Analysis and Processing. (1995)
13. Canny, J.: A computational approach to edge detection. IEEE Trans. Pattern Anal. Mach. Intell. **8** (1986) 679–698