Motion Field of Curves: Applications*

Théo Papadopoulo and Olivier Faugeras

INRIA Sophia Antipolis,
2004 Route des Lucioles,
BP 93, 06902 SOPHIA-ANTIPOLIS Cedex, FRANCE
papadop@sophia.inria.fr, faugeras@sophia.inria.fr

Abstract. This paper discusses the well known problem of structure from motion for the special case of rigid curves. It is already known that it is theoretically possible to recover the motion and thus the structure of a moving 3D rigid curve observed through one camera given some set of derivatives that are defined on the so-called spatio-temporal surface under the most general camera model of perspective projection. We give here a new simplification of the previous results. In order to show that implementing this theory is indeed feasible, we proceeded towards two main directions. First, we have implemented the special case of planar rigid curves. Second, we show that the derivatives defined on the spatio-temporal surface which are needed in the general case can indeed be computed from the images.

1 Introduction

Recovering three-dimensionnal information about an observed scene from images is one of the main goals of computer vision. The basic underlying idea leading to a vast majority of methods is to combine information coming from many different viewpoints.

One way to obtain multiple viewpoints is to use motion [FLT87, WKPS87, FDN89, SA90]. In these approaches, only one camera is used. From the measurement of the motion of image primitives (flow fields), the 3D motion parameters and the relative depth can be computed. However, there are many ways to compute flow fields: most of these approaches are based on optical flow which is computed directly from image intensities [LHP80, HS81, Nag83, Hil84, Koe86, D'H86, Bou89, Gon89]. Another way to proceed is to use the motion fields in the image which are defined as the projection in the image of the 3D motion field of some geometric object (usually points or lines). We have studied in [FP93] the relationship between optical flow and motion field for general 3D curves and shown that the assumptions usually made in the computation of the optical flow are a bit difficult to defend. Moreover, we have shown that for a rigid 3D curve it is theoretically possible to recover the 3D structure and motion from a monocular sequence of images. The data needed in order to achieve these tasks are derivatives defined on the so-called spatio-temporal surface. In order to implement these ideas, we focussed on three main directions:

- Simplify the existing equations.

* The research described in this paper has been supported by DRET contract N° 911 349/A00 and by ESPRIT BRA project Insight II

Lecture Notes in Computer Science, Vol. 800
Jan-Olof Eklundh (Ed.)
Computer Vision – ECCV '94

- Find methods to compute the needed derivatives.
- Implement some simpler cases, namely the case of planar 3D curves.

After the introduction of notations, we examine briefly why the equations for the general case can be simplified and how to relate the 3D motion of planar curves to the first order derivatives in space and time on the spatio-temporal surface. In the final section, we discuss our implementation for the planar curve case and results on both synthetic and real images. As well, results of the derivative computation needed for the general monocular case are presented.

2 Notations and Basic Results

The goal of this section is to introduce some notations relative to the problem of recovering the 3D motion from 2D motion fields. We also recall some basic results that are described in details in [FP93] and explain without proof why one of the basic results of this paper can be simplified.

2.1 The Camera Model

We assume that the camera obeys the standard pinhole model with unit focal length. We note O the focal center point and suppose that the retina $\mathcal{R}$ is parallel to the plane (O, X, Y). The frame (O, X, Y, Z) being naturally attached to the camera model, all equations involving 3D parameters are written in this frame.

Given a 3D point $\mathbf{M} = (X, Y, Z)$ and its 2D perspective projection $\mathbf{m} = (x, y, 1)$ on the $\mathcal{R}$ plane, their relationship is characterized by the following equation:

$$\mathbf{M} = Z\mathbf{m} \tag{1}$$

This equation is fundamental in that all the constraints we present here are direct consequences of it. The concept of temporal variation can be incorporated with the introduction of a time factor τ.

2.2 Definitions

We now assume that we observe in a sequence of images a family (c_τ) of curves, where τ denotes time, which we assume to be the perspective projection in the retina of a 3D curve (C) that moves in space. If we consider the three-dimensional space (x, y, τ), this family of curves sweeps in that space a surface (Σ) defined as the set of points $((c_\tau), \tau)$. Figure 1 illustrates an example of one such spatio-temporal surface generated by a circle rotating around one of its diameters in front of the camera.

At a given time instant τ, let s be the arclength of (c_τ) and S the arclength of (C). We further suppose that S is not a function of time (i.e. the motion is isometric). Now, for a point $\mathbf{m}$ on (c_τ), it is possible to define two different motion fieds: the *apparent motion field* $\mathbf{v}_m^a$ and the *real motion field* $\mathbf{v}_m^r$ of $\mathbf{m}(s, \tau)$ obtained by taking the partial derivative of $\mathbf{m}(s, \tau)$ with respect to time when respectively s or S is kept constant.

Introducing the Frenet frame $(\mathbf{t}, \mathbf{n})$, where $\mathbf{n}$ is the unit normal vector to (c_τ) at $\mathbf{m}$, and under the weak assumption of *isometric* motion, we reach the following conclusions from the study of the spatio-temporal surface:

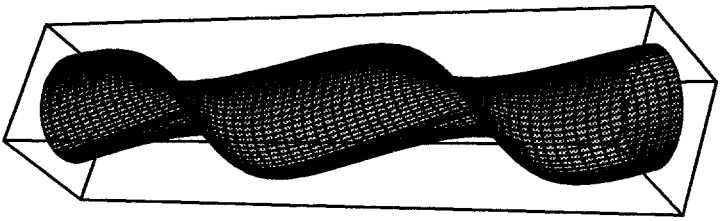

Fig. 1. The spatio-temporal surface generated by a circle rotating in front of the camera

1. The normal motion field β is the same for $\mathbf{v}_m^a$ and $\mathbf{v}_m^r$ and can be recovered from the normal to the spatio-temporal surface,
2. the tangential apparent motion field α can be recovered from the normal motion field,
3. the tangential real motion field w *cannot* be recovered from the spatio-temporal surface.

Therefore, the full real motion field is not computable from the observation of the image of a moving curve under the isometric assumption. This can be considered as a new statement of the so-called *aperture* problem. In order to solve it we *must* add more constraints, for example that the 3D motion is rigid.

2.3 The Case of a Rigid 3D Curve

Assuming now that the curve (C) is moving rigidly. Let $(\Omega, \mathbf{V})$ be its kinematic screw at the optical center O of the camera. We assume also that the camera has been normalized by calibration to unit focal length.

Taking the total derivative of (1) with respect to time, using the standard formula giving the velocity $\dot{\mathbf{M}}$ of any point $\mathbf{M}$ attached to the rigid body as a function of the kinematic screw and projecting this vector equation onto $\mathbf{t}$ and $\mathbf{n}$ yields two scalar equations:

$$Z(w + \Omega \cdot \mathbf{b}) = -\mathbf{U_n} \cdot \mathbf{V} \tag{2}$$

$$Z(\beta - \Omega \cdot \mathbf{a}) = \mathbf{U_t} \cdot \mathbf{V} \tag{3}$$

where $\mathbf{U_t}$, $\mathbf{U_n}$, $\mathbf{a}$, $\mathbf{b}$ denote respectively $\mathbf{m} \times \mathbf{t}$, $\mathbf{m} \times \mathbf{n}$, $\mathbf{m} \times \mathbf{U_t}$ and $\mathbf{m} \times \mathbf{U_n}$.

These equations are fundamental (especially (3)) in the sense that they express the relationship between the unknown 3D motion of a point and the real motion field of its image.

Let's now recall some previous result that is given in [FP93].

Theorem 1. *At each point of an observed curve (c_τ) evolving during time, it is possible to write two polynomial equations in the coordinates Ω, $\mathbf{V}$, $\dot{\Omega}$ and $\dot{\mathbf{V}}$ (The kinematic*

screw attached to the 3D curve and its first time derivative). The coefficients of these equations are polynomials in the quantities

$$\begin{matrix} \beta & \frac{\partial \beta}{\partial s} & \frac{\partial^2 \beta}{\partial s^2} & \partial_{\mathbf{n}_\beta} \beta & \partial_{\mathbf{n}_\beta} \frac{\partial \beta}{\partial s} \\ \kappa & \frac{\partial \kappa}{\partial s} & & \partial_{\mathbf{n}_\beta} \kappa & \end{matrix}$$

that can be measured from the spatio-temporal surface (Σ).

These equations lead to a way to compute the motion and the structure of the 3D curve. It can be shown that they are, in fact, a direct consequence of (3) and that one of these equations is redundant because it can be expressed as a linear combination of the other equation and its first order time derivative. Thus the previous result can be restated as:

Theorem 2. *At each point of an observed curve* (c_τ) *evolving during time, it is possible to write one polynomial equation in the coordinates* Ω, $\mathbf{V}$, $\dot{\Omega}$ *and* $\dot{\mathbf{V}}$ *(The kinematic screw attached to the 3D curve and its first time derivative). The coefficients of this equation are polynomials in the quantities*

$$\begin{matrix} \beta & \frac{\partial \beta}{\partial s} & \partial_{\mathbf{n}_\beta} \beta \\ \kappa & & \end{matrix}$$

that can be measured from the spatio-temporal surface (Σ).

The nice thing with this new theorem is that we get rid of all third order derivatives with only first and second order derivatives being left. Therefore we are only interested in the above-mentioned derivatives and we show later that it is possible to compute them quite precisely.

3 The Motion of 3D Planar Rigid Curves

We study here a special case of the motion of rigid curves sketched in the previous section: the case of a 3D rigid curve that is planar. By making this hypothesis, it is possible to write an equation similar to those obtained in the general case but:

- in which $\dot{\Omega}$ and $\dot{\mathbf{V}}$ are no longer involved, thus leading to a system of equations with less unknowns.
- of total degree 2 (instead of 4).
- in which only the first order derivatives of Σ appear.

The first two properties show that the systems we obtain is much simpler (less unknowns with lower degrees), thus the number of possible solutions is smaller. The second characteristic means that not only the equation are simpler but also they are more stable with respect to the measurement noise. We thus may hope that the solutions of the system are also more stable.

3.1 The Equation in the Planar Case

Let $\mathbf{M} = [X, Y, Z]^T$ be a point on the 3D planar curve. This point belongs to the plane of the curve. If we suppose that the image of the curve is not degenerated into a segment, then the optical center of the camera is not on the curve plane. The point M thus verifies the following equation:

$$\mathbf{N} \cdot \mathbf{M} + 1 = aX + bY + cZ + 1 = 0 \tag{4}$$

Combining (4) with the perspective equation (1), we obtain:

$$Z = \frac{-1}{\mathbf{N} \cdot \mathbf{m}} \tag{5}$$

Equation (5) is fundamental because it connects the plane structure of the curve Z to the measures in the image x, y. Practically, it allows to replace the quantity Z that varies along the observed curve by 3 quantities a, b, c that are constant along this same curve. Replacing Z by its value given by this equation in (3), yields the following theorem:

Theorem 3. *At each point of* (c_r) *considered as the projection of a 3D planar curve, it is possible to write a polynomial equation in the unknowns* Ω, $\mathbf{V}$, $\mathbf{N}$.

$$\beta - \Omega \cdot \mathbf{a} + (\mathbf{U_t} \cdot \mathbf{V})(\mathbf{N} \cdot \mathbf{m}) = 0 \tag{6}$$

This equation is not homogeneous in $\mathbf{V}$ *but we can see that if* $(\mathbf{V}, \mathbf{N})$ *is a solution then* $(\lambda\mathbf{V}, \frac{1}{\lambda}\mathbf{N})$ *is also a solution for every* $\lambda \neq 0$. *This property shows that, as in the general case, only the direction of* $\mathbf{V}$ *can be recovered. The equation is of degree 2 in* $(\mathbf{V}, \mathbf{N})$, *of degree 1 in* Ω *and of total degree 2.*

In the following, we call this equation the planar equation. Evaluating this equation at 8 points, we obtain a system of degree 2 in the 9 unknowns $(\Omega, \mathbf{V}, \mathbf{N})$. It is then possible to reformulate the conjecture we have made in the previous section for this particular case:

Conjecture 4. *The kinematic screw* $(\Omega, \mathbf{V})$ *and the normal* $\mathbf{N}$ *to the plane of a 3D rigid planar curve can, in general, be estimated from the observation of the spatio-temporal surface generated by its image on the retina by solving a system of polynomial equations. The quantity* Z *can be estimated at each point up to a scale factor by using (5).*

Of course, as in the general case, this conjecture is wrong in some special cases such as straight lines or conics. See [Ber89] for other examples of ambiguity. Practically, for non ambiguous curves, this conjecture has always been found to be true.

3.2 Reconstruction of the Curve

From previous formulas, once the motion is computed, it is possible to reconstruct the 3D curve up to a scale factor by two different means:

- Using (4) which relates Z to the plane parameters $\mathbf{N}$. Here we are using explicitely the planar hypothesis.
- Using the general (3) that is true for all 3D rigid curve. It relates Z to the kinematic screw $(\Omega, \mathbf{V})$. Since we do not use the planar hypothesis, the reconstructions computed this way are more unstable than the previous ones.

3.3 Ambiguity of the Solutions

We are interested here in describing the structure of the solutions. The question is: is it possible generically to find a relation between two solutions (here this means that the result we look for does not depend on the actual values of the estimated parameters). One way to do this is to search for a transformation on $(\Omega, \mathbf{V}, \mathbf{N})$ that leaves the equation unchanged.

Let us thus look at the coefficients of (6) in variables x, y, $\mathbf{t}$, β. These different terms represent the way the information relative to $(\Omega, \mathbf{V}, \mathbf{N})$ is coded in (6) and this is independent of the point at which the equation is written. In some way, every quantity that cannot be computed from these terms or that remains ambiguous exhibits the same behaviour when computed from (6). We use this property to prove that there is a companion solution to each solution of the system.

Theorem 5. *If* $(\Omega, \mathbf{V}, \mathbf{N})$ *is a solution of the system obtained for a planar curve then* $(\Omega + \mathbf{V} \wedge \mathbf{N}, \mathbf{N}, \mathbf{V})$ *is also a solution of this same system.*

This theorem is the specialization to planar curves of a well-known theorem on planar points [TH82, LH84, May92]. This is not surprising since a planar curve is nothing more than a set of planar points. What is more surprising however is that it can be shown that there is no new ambiguity introduced by the fact we only use normal flow information.

The coefficients of (6) in the variables x, y, $\mathbf{t}$, β can also be used to show that there is at least one and at most 3 solutions to this problem (counting only once the two related solutions).

4 Implementation

We describe here the implementation of the theory described in the previous section and show the results we have obtained on both synthetic and real images.

4.1 Motion of Planar Curves

Many tests have shown that it is difficult to estimate with a good accuracy spatio-temporal parameters. Temporal derivatives are especially difficult to obtain: this phenomenon seems to come from sampling problems in time. Whereas spatial sampling of an image may be known and constant (it is fixed by the physical parameters of the camera), time sampling of the spatio-temporal surface Σ around a point P should depend on the speed of that point. Temporal derivatives can be obtained very easily by considering the curve drawn on Σ that lies in the plane defined by the point $\mathbf{m}$ at which we want to compute the spatio-temporal parameters, and being spanned by the vectors $\mathbf{n}$ the normal vector to the observed curve at $\mathbf{m}$ and by τ the unit vector on the time axis. Therefore, the accuracy of the temporal derivatives depends upon how well this curve is sampled.

To validate the approach, we use two image sequences (of about 30 to 40 images each). The first one (see Fig. 2 left) is a synthetic sequence of a planar 3D quartic. The second one (see Fig. 2 right) is a real sequence. In all these images the 3D curve rotates around a vertical axis and translates in the same direction.

Here is the general scheme of the implementation:

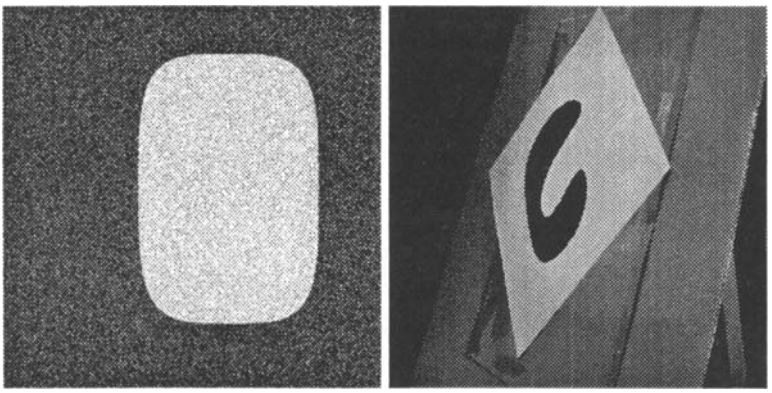

Fig. 2. Left: an image exerpted from the synthetic sequence. A Gaussian noise of signal/noise ratio of 20% has been added on intensities. Right: an image exerpted from the real sequence.

- Build the spatio-temporal surface.
- Estimate the spatio-temporal parameters at every point of the curve at one time instant.
- Normalise the parameters and finally solving of the polynomial system.

Each of these stages is described in the next paragraphs.

4.2 Building the Spatio-Temporal Surface

In order to speed up computation times, we gather the points of the spatio-temporal surface in a data structure that allows us to compute easily the neighbours of a given point on the surface. The solution adopted is to link the points together using two doubly linked list: one for spatial neighbors and the other for time neighbors. The algorithm attempts to use the continuity of the curves to avoid walking through too many points of either (c_τ) or $(c_{\tau+d\tau})$.

4.3 Estimating Spatio-Temporal Parameters

The estimation of the spatio-temporal parameters is just sketched here. Because of the discrepancy between space and time sampling rates we compute independently spatial and temporal parameters. First, the local orientation at each point is computed: to do so, we construct the two signals $x(s)$ and $y(s)$ and fit locally models to them. Deriving these models gives the local derivatives $x'(s)$ and $y'(s)$ that describe the local tangent. Thus, the angle $\theta(s)$ between the normal and the horizontal is obtained.

The left part of Fig. 3 shows the angle estimates along the curve. The maximal error between the theoretical curve and the measures is 0.011 radians.

We then compute the value of the β parameter at each point. At each point $\mathbf{m}$, which orientation is given by $\mathbf{n}$, we build the curve that is defined as the intersection of the spatio-temporal surface Σ with the plane Π defined by $\mathbf{m}$, $\mathbf{n}$ and τ the time axis. This curve can be represented in the plane (τ, d) where d is the distance in the direction

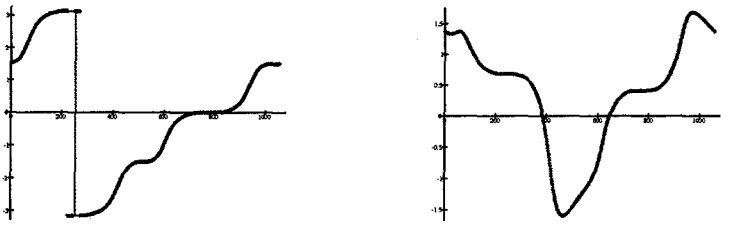

Fig. 3. Estimated angle (left) and β (right) along the curve. The X-axis is curvilinear abscissa. The crosses represent measures whereas plain curves represent theoretical values of the angle.

of n between a point of plane Π and m (this distance is also the distance of this point to the tangent to (c_τ) at m if we represent (c_τ) and $(c_{\tau+d\tau})$ in the same plane). We thus obtain a curve that is approximated by a truncated Chebyshev polynomial (to reduce the effects of noise). The derivative of this polynomial at time τ is then computed and is nothing but β.

The right part of Fig. 3 shows the β estimates along the curve at one time instant for the synthetic image sequence. The maximal error between the theoretical curve and the measures is 0.017152 pixels by image (the image number here is the time coordinate).

4.4 Parameters Normalisation and System Resolution

At this point, we compute the normalized values of all the spatio-temporal parameters we need. Then we build the polynomial system obtained from the plane equation expressed at each of these points. We arbitrarily normalize the first component of N to 1 and obtain a system in 7 unknowns and as many equations as there are points on the edge. We then solve this system using a modified Newton method.

For the synthetic sequence, using the theoretical values of the parameters we have mathematically proved that there are only two solutions to the motion problem. Namely the true solution and its companion solution as described in Theorem 5. In what follows we always consider errors between the true solution and the corresponding solution.

In Table 1, we have tabulated the accuracy of the computed solution for all the components of (Ω, V, N) as a function of the number of images of the synthetic sequence (around image 14) used to compute the spatio-temporal parameters. Note that (and this is be true for all results showed here) for V and N only the angle between the theoretical and estimated values are shown since these vectors are only defined up to a scale factor. The time needed to do all the computations (including parameter estimation and resolution of the system which has 1112 equations) is about 32 seconds when 5 images are used and about 36 seconds when all 29 images are used. These times have been obtained on a Sun Sparc 2.

Figure 4 shows the values of Z along the curve and the reconstruction for the good solution. Note that there are always points on the curves where the value of Z computed from (3) is biased: it can be shown that these points corresponds to the

Number of images considered	$\| \Omega - \Omega_{\text{theor}} \|$ in °/image	$\widehat{\Omega, \Omega}_{\text{theor}}$ in °	$\widehat{\mathbf{V}, \mathbf{V}}_{\text{theor}}$ in °	$\widehat{\mathbf{N}, \mathbf{N}}_{\text{theor}}$ in °
5	1.6×10^{-1}	4.1×10^{-1}	1.2	1.2
13	2.4×10^{-2}	1.1	1.0×10^{-1}	6.0×10^{-2}
20	1.4×10^{-2}	6.4×10^{-1}	6.4×10^{-1}	6.2×10^{-2}
29	1.1×10^{-2}	5.4×10^{-1}	5.1×10^{-1}	5.6×10^{-2}

Table 1. Errors in norm and angle between the estimated results and the theoretical ones as a function of the number of images used to compute the spatio-temporal parameters.

points for which $\mathbf{U_t} \cdot \mathbf{V} = 0$ which can be interpreted geometrically as the points at which the tangent to the curve goes through the focus of expansion. As it can be shown from (3), there is no depth information at these points.

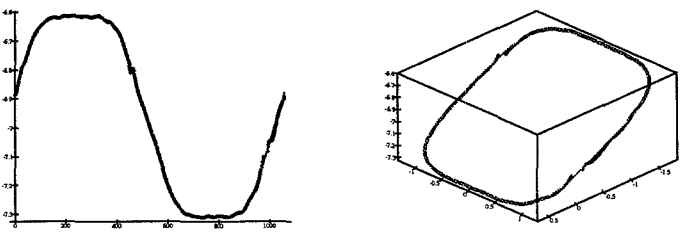

Fig. 4. On the left: estimates of Z along the curve. On the right: 3D reconstructions for the correct solution. The plain curve is the reconstruction based on (4) whereas the crosses represent the reconstruction based on (3).

A nice experimental consequence resulting from the comparison of the planar and general reconstruction is that it seems possible to distinguish automatically the correct solution from the bad one: the correct solution is always associated to the reconstructions for which the errors are the smallest.

Verifying quantatively the results obtained with the real sequences is a difficult task: the best way to do it is to look at the angular speed. With the measured data of Fig. 5 the measured angular speed is $-1.055°$/image where it should be $-1°$/image. For this sequence, the worst relative error on the angular speed is under 15% but usually this relative error is lower than 7%. Note also that we have a good robustness to wrong estimation of camera parameters: we used many different internal parameters (involving changes of 5 to 10 pixels for the optical center and changes of 2 to 5% for the scale factors) obtained by calibrating with different data and noticed a good stability of the computed motion (the relative error between the different parameters is at most 1%).

Figure 5 shows the angle and beta estimates along the curve. Figure 6 shows the reconstructions obtained from the correct and the bad solutions.

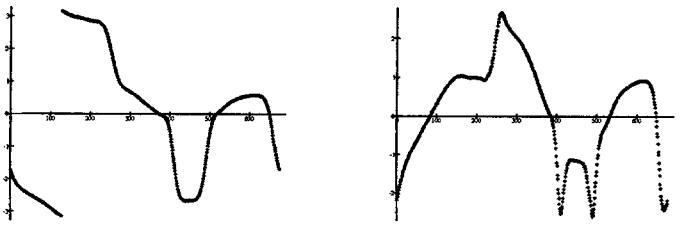

Fig. 5. Estimated angle (left) and β (right) along the curve. The X-axis is curvilinear abscissa. The crosses represent measures whereas plain curves represent theoretical values of the angle.

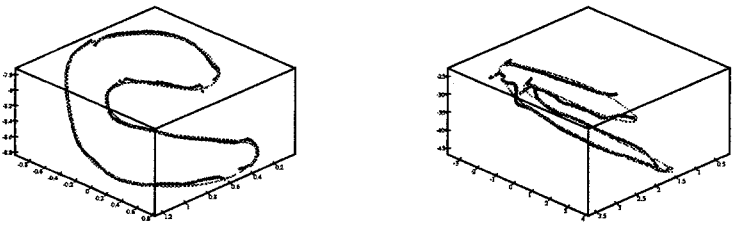

Fig. 6. 3D reconstructions for the correct solution (left) and the bad solution (right). The plain curve is the reconstruction based on (4) whereas the crosses represent the reconstruction based on (3).

4.5 Higher Order Derivatives

We conclude by showing some figures giving the parameters κ, $\frac{\partial \beta}{\partial s}$ and $\partial_{\mathbf{n}_\beta}\beta$ along the curve for the synthetic image sequence. These measures seem good enough to allow a practical implementation of the stereo disambiguation described in [FP93] as well as that of the general 3D rigid curve case based on Theorem 2.

Figure 7 shows the parameters κ, $\frac{\partial \beta}{\partial s}$ and $\partial_{\mathbf{n}_\beta}\beta$ along the observed curve. There are still some problems around curvature extrema. In fact, the origin of these problems is now well understood and will be corrected in future work.

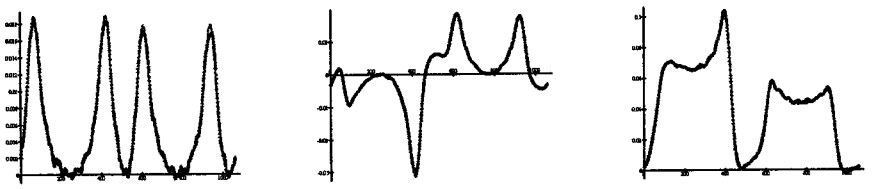

Fig. 7. From left to right: measured $\kappa, \frac{\partial \beta}{\partial s}$ and $\partial_{\mathbf{n}_\beta} \beta$ along the curve. Crosses represent measures whereas the plain curve represent the theoretical values.

5 Conclusion and Future Work

We have shown in this paper how one of the two equations relating the 3D motion of a rigid curve to the observed motion field is redundant. Moreover, we have shown that the remaining equation does not depend on the third order derivatives on the spatio-temporal surface generated by the retinal image of the 3D curve and demonstrated that the first and second order derivatives can indeed be computed quite accurately from the long image sequences. This is very promising for the implementation of that theory which is, obviously, our next goal.

From another point of view, we have completely implemented the more specific case of planar rigid 3D curves and shown on real and synthetic images that the theory gives quite accurate results. Moreover, it seems that the fundamental ambiguity related to the use of this special kind of curves can be overcome by looking at the errors on the Z reconstruction errors between the planar and general model. This case was most fruitful since it allowed us to understand better many of the characteristics of the systems of polynomials and to develop many tools that are useful for implementing the general case. Many improvements can still be made to improve further the quality of the results such as working with many planar curve patches or taking into account special properties of some points such as inflexion points, bitangent points or points at which the tangent to the curve goes through the focus of expansion.

References

[Ber89] Fredrik Bergholm. Motion from Flow along Contours: A Note on Robustness and Ambiguous Cases. *The International Journal of Computer Vision*, 2(4):395–415, April 1989.

[Bou89] Patrick Bouthemy. A Maximum Likelihood Framework for Determining Moving Edges. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(5):499–511, May 1989.

[D'H86] Johan D'Hayer. Determining Motion of Image Curves from Local Pattern Changes. *Computer Vision, Graphics, and Image Processing*, 34:166–188, 1986.

[FDN89] Olivier D. Faugeras, Nourr-Eddine Deriche, and Nassir Navab. From optical flow of lines to 3D motion and structure. In *Proceedings IEEE RSJ International Workshop on Intelligent Robots and Systems '89*, pages 646–649, 1989. Tsukuba, Japan.

[FLT87] Olivier D. Faugeras, Francis Lustman, and Giorgio Toscani. Motion and Structure from point and line matches. In *Proceedings of the First International Conference on Computer Vision, London*, pages 25–34, June 1987.

[FP93] Olivier D. Faugeras and Théo Papadopoulo. A theory of the motion fields of curves. *The International Journal of Computer Vision*, 10(2):125–156, 1993.

[Gon89] S. Gong. Curve Motion Constraint Equation and its Applications. In *Proceedings Workshop on Visual Motion*, pages 73–80, 1989. Irvine, California, USA.

[Hil84] Ellen C. Hildreth. *The Measurement of Visual Motion.* MIT Press, Cambridge, Mass., 1984.

[HS81] Berthold K. P. Horn and Brian G. Schunk. Determining Optical Flow. *Artificial Intelligence*, 17:185–203, 1981.

[Koe86] Jan J. Koenderink. Optic Flow. *Vision Research*, 26(1):161–180, 1986.

[LH84] H.C. Longuet-Higgins. The visual ambiguity of a moving plane. *Proceedings of the Royal Society London, B*, 223:165–175, 1984.

[LHP80] H. C. Longuet-Higgins and K. Prazdny. The interpretation of moving retinal images. *Proceedings of the Royal Society of London*, B 208:385–387, 1980.

[May92] S.J. Maybank. *Theory of reconstruction From Image Motion.* Springer-Verlag, 1992.

[Nag83] H-H. Nagel. Displacement Vectors Derived from Second Order Intensity Variations in Image Sequences. *Computer Vision, Graphics, and Image Processing*, 21:85–117, 1983.

[SA90] Minas E. Spetsakis and John Aloimonos. Structure from Motion Using Line Correspondences. *The International Journal of Computer Vision*, 4:171–183, 1990.

[TH82] R. Tsai and T.S. Huang. Estimating three-dimensional motion parameters of a rigid planar patch, ii: singular value decomposition. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 30, 1982.

[WKPS87] Allen M. Waxman, Behrooz Kamgar-Parsi, and Muralidhara Subbarao. Closed-Form Solutions to Image Flow Equations for 3D Structure and Motion. *The International Journal of Computer Vision*, 1:239–258, 1987.