Reconstruction of 3D Curves for Quality Control

Hanna Martinsson¹, Francois Gaspard², Adrien Bartoli²,
and Jean-Marc Lavest²

¹ CEA, LIST, Boîte Courrier 94, F-91 191 Gif sur Yvette, France
`hanna.martinsson@cea.fr`

² LASMEA (CNRS/UBP), 24 avenue des Landais, F-63 177 Aubière, France
`adrien.bartoli@gmail.fr`

Abstract. In the area of quality control by vision, the reconstruction of 3D curves is a convenient tool to detect and quantify possible anomalies. Whereas other methods exist that allow us to describe surface elements, the contour approach will prove to be useful to reconstruct the object close to discontinuities, such as holes or edges.

We present an algorithm for the reconstruction of 3D parametric curves, based on a fixed complexity model, embedded in an iterative framework of control point insertion. The successive increase of degrees of freedom provides for a good precision while avoiding to over-parameterize the model. The curve is reconstructed by adapting the projections of a 3D NURBS snake to the observed curves in a multi-view setting.

1 Introduction

The use of optical sensors in metrology applications is a complicated task when dealing with complex or irregular structures. More precisely, projection of structured light allows for an accurate reconstruction of surface points but does not allow for a precise localization of the discontinuities of the object. This paper deals with the problem of reconstruction of 3D curves, given the CAD model, for the purpose of a control of conformity with respect to this model. We dispose of a set of images with given perspective projection matrices. The reconstruction will be accomplished by means of the observed contours and their matching, both across the images and to the model.

Algorithms based on active contours [9] allows for a local adjustment of the model and a precise reconstruction of primitives. More precisely, the method allows for an evolution of the reprojected model curves toward the image edges, thus to minimize the distance in the images between the predicted curves and the observed edges.

The parameterization of the curves as well as the optimization algorithms we use must yield an estimate that meets the requirements of accuracy and robustness necessary to perform a control of conformity. We have chosen to use NURBS curves [11], a powerful mathematical tool that is also widely used in industrial applications.

In order to ensure stability, any method used ought to be robust to erroneous data, namely the primitives extracted from the images, since images of metallic objects incorporate numerous false edges due to reflections.

Although initially defined for ordered point clouds, active contours have been adapted to parametric curves. Cham and Cipolla propose a method based on affine epipolar geometry [3] that reconstructs a parametric curve in a canonical frame using stereo vision. The result is two coupled snakes, but without directly expressing the 3D points. In [15], Xiao and Li deal with the problem of reconstruction of 3D curves from two images. However, the NURBS curves are approximated by B-splines, which makes the problem linear, at the expense of losing projective invariance. The reconstruction is based on a matching process using epipolar geometry followed by triangulation. The estimation of the curves is performed independently in the two images, that is, there is no interactivity between the 2D observations and the 3D curve in the optimization. Kahl and August introduce in [8] a coupling between matching and reconstruction, based on an a priori distribution of the curves and on an image formation model. The curves are expressed as B-splines and the optimization is done using gradient descent.

Other problems related to the estimation of parametric structures have come up in the area of surfaces. In [14], Siddiqui and Sclaroff present a method to reconstruct rational B-spline surfaces. Point correspondences are supposed given. In a first step, B-spline surface patches are estimated in each view, then the surface in 3D, together with the projection matrices, are computed using factorization. Finally, the surface and the projection matrices are refined iteratively by minimizing the 2D residual error. So as to avoid problems due to over-parameterization, the number of control points is limited initially, to be increased later on in a hierarchical process by control point insertion.

In the case of 2D curve estimation, other aspects of the problem are addressed. Cham and Cipolla adjust a spline curve to fit an image contour [4]. Control points are inserted iteratively using a new method called PERM (potential for energy-reduction maximization). An MDL (minimal description length [7]) strategy is used to define a stopping criterion. In order to update the curve, the actual curve is sampled and a line-search is performed in the image to localize the target shape. The optimization is performed by gradient descent. Brigger et al. present in [2] a B-spline snake method without internal energy, due to the intrinsic regularity of B-spline curves. The optimization is done on the knot points rather than on the control points, which allows the formulation of a system of equations that can be solved by digital filtering. So as to increase numerical stability, the method is embedded in a multi-resolution framework. Meegama and Rajapakse introduce in [10] an adaptive procedure for control point insertion and deletion, based on the euclidean distance between consecutive control points and on the curvature of the NURBS curve. Local control is ensured by adjustment of the weights. The control points evolve in each iteration in a small neighborhood (3×3 pixels). Yang et al. use a distance field computed a priori with the fast marching method in order to adjust a B-spline snake [16]. Control points are

added in the segment presenting a large estimation error, due to a degree of freedom insufficient for a good fit of the curve. The procedure is repeated until the error is lower than a fixed threshold. Redundant control points are then removed, as long as the error remains lower than the threshold.

2 Problem Formulation

Given a set of images of an object, together with its CAD model, our goal is to reconstruct in 3D the curves observed in the images. In order to obtain a 3D curve that meets our requirements regarding regularity, rather than reconstructing a point cloud, we estimate a NURBS curve. The reconstruction is performed by minimizing the quadratic approximation error. The minimization problem is formulated for a set of M images and N sample points by

$$\mathbf{C}(\mathcal{P}) = \arg \min_{\mathcal{P}} \sum_{i=0}^{M-1} \sum_{j=0}^{N-1} (\mathbf{q}_{ij} - T_i(\mathbf{C}(\mathcal{P}, t_j)))^2, \quad (1)$$

where $\mathbf{q}_{ij}$ is a contour point associated with the curve point of parameter t_j , T_i is the projective operator for image i and $\mathcal{P}$ is the set of control points.

Our choice to use NURBS curves is justified by several reasons. First, NURBS curves have interesting geometrical properties, namely concerning regularity and continuity. An important geometrical property that will be of particular interest is the invariance under projective transformations.

3 Properties of NURBS Curves

Let $U = \{u_0, \dots, u_m\}$ be an increasing vector, called the knot vector. A NURBS curve is a vector valued, piecewise rational polynomial over U , defined by

$$\mathbf{C}(t) = \sum_{i=0}^n \mathbf{P}_i R_{i,k}(t) \quad \text{with} \quad R_{i,k}(t) = \frac{w_i B_{i,k}(t)}{\sum_{j=0}^n w_j B_{j,k}(t)}, \quad (2)$$

where $\mathbf{P}_i$ are the control points, $B_{i,k}(t)$ the B-spline basis functions defined over U , w_i the associated weights and k the degree.

It is a common choice to take $k = 3$, which has proved to be a good compromise between required smoothness and the problem of oscillation, inherent to high degree polynomials. For our purposes, the parameterization of closed curves, we consider periodic knot vectors, that is, verifying $u_{j+m} = u_j$. Given all these parameters, the set of NURBS defined on U forms, together with the operations of point-wise addition and multiplication with a scalar, a vector space.

For details on NURBS curves and their properties, refer to [11].

3.1 Projective Invariance

According to the pinhole camera model, the perspective projection $T(\cdot)$ that transforms a world point into an image point is expressed in homogeneous coordinates by means of the transformation matrix $\mathbf{T}_{3 \times 4}$. Using weights associated

with the control points, NURBS curves have the important property of being invariant under projective transformations. Indeed, the projection of (2) remains a NURBS, defined by its projected control points and their modified weights. The curve is written

$$\mathbf{c}(t) = T(\mathbf{C})(t) = \frac{\sum_{i=0}^n w'_i T(\mathbf{P}_i) B_{i,k}(t)}{\sum_{i=0}^n w'_i B_{i,k}(t)} = \sum_{i=0}^n T(\mathbf{P}_i) R'_{i,k}(t) \quad (3)$$

The new weights are given by

$$w'_i = (T_{3,1}X_i + T_{3,2}Y_i + T_{3,3}Z_i + T_{3,4}) w_i = \mathbf{n} \cdot (\mathbf{C}_O - \mathbf{P}_i) w_i, \quad (4)$$

where $\mathbf{n}$ is a unit vector along the optical axis and $\mathbf{C}_O$ the optical center of the camera.

3.2 Control Point Insertion

One of the fundamental geometric algorithms available for NURBS is the control point insertion. The key is the knot insertion, which is equivalent to adding one dimension to the vector space, consequently adapting the basis. Since the original vector space is included in the new one, there is a set of control points such that the curve remains unchanged.

Let $\bar{u} \in [u_j, u_{j+1})$. We insert $\bar{u}$ in U , forming the new knot vector $\bar{U} = \{\bar{u}_0 = u_0, \dots, \bar{u}_j = u_j, \bar{u}_{j+1} = \bar{u}, \bar{u}_{j+2} = u_{j+1}, \dots, \bar{u}_{m+1} = u_m\}$. The new control points $\bar{\mathbf{P}}_i$ are given by the linear system

$$\sum_{i=0}^n \mathbf{P}_i R_{i,k}(t) = \sum_{i=0}^{n+1} \bar{\mathbf{P}}_i \bar{R}_{i,k}(t). \quad (5)$$

We present the solution without proof. The new control points are written

$$\bar{\mathbf{P}}_i = \alpha_i \mathbf{P}_i + (1 - \alpha_i) \mathbf{P}_{i-1}, \quad (6)$$

with

$$\alpha_i = \begin{cases} 1 & i \leq j - k \\ \frac{\bar{u} - u_i}{u_{i+k} - u_i} & \text{if } j - k + 1 \leq i \leq j \\ 0 & i \geq j + 1 \end{cases}. \quad (7)$$

Note that only k new control points need to be computed, due to the local influence of splines.

4 Curve Estimation

Using the NURBS formulation, the minimization problem (1) is written

$$\min_{\{\hat{\mathbf{P}}_l\}} \sum_{i=0}^{M-1} \sum_{j=0}^{N-1} \left(\mathbf{q}_{ij} - \sum_{l=0}^n T_i(\hat{\mathbf{P}}_l) R_{l,k}^{(i)}(t_j) \right)^2, \quad (8)$$

where $R_{l,k}^{(i)}$ are the basis functions for the projected NURBS curve in image i .

The problem has two parts. First, the search for candidate edge points $\mathbf{q}_{ij}$ in the images, then the optimization of the 3D NURBS curve by optimization on the control points. The search for candidate points is carried out independently in the images using a method inspired by the one used by Drummond and Cipolla in [6]. For the optimization problem, we use the non-linear Levenberg-Marquardt minimization method. This optimization allows the control points to move in 3D, but does not change their number. In order to obtain an optimal reconstruction of the observed curve, we iteratively perform control point insertion.

4.1 Search for Image Contours

We sample the NURBS curve projected in the image, to use as starting points in the search for matching contour points. A line-search is performed in order to find the new position of the curve, ideally corresponding to an edge. Our approach is based solely on the contours. Due to the aperture problem, the component of motion of an edge, tangent to itself, is not detectable locally and we therefore restrict the search for the new edge position to the edge normal at each sample point. As we expect the motion to be small, we define a search range (typically in the order of 20 pixels) so as to limit computational cost. In order to find the new position of a sample point, for each point belonging to the normal within the range, we evaluate the gradient and compute a weight based on the intensity and the orientation of the gradient and the distance from the sample point. The weight function v_j for a sample point p_j and the candidate point p_ξ will be of the form

$$v_j(p_\xi) = \varphi_1(|\nabla I_\xi|) \cdot \varphi_2\left(\frac{\hat{n}_j \cdot \nabla I_\xi}{|\nabla I_\xi|}\right) \cdot \varphi_3(|p_j - p_\xi|),$$

where $\hat{n}_j$ is the normal of the projected curve at sample point j , ∇I_ξ is the gradient at the candidate point and the φ_k are functions to define. The weight function will be evaluated for each candidate p_ξ and the point p'_j with the highest weight, identified by its distance from the original point $d_j = |p_j - p'_j|$, will be retained as the candidate for the new position of the point.

The bounded search range and the weighting of the point based on their distance from the curve yield a robust behavior, close to that of an M-estimator.

4.2 Optimization on the Control Points

The first step of the optimization consists in projecting the curve in the images. Since the surface model is known, we can identify the visible parts of the curve in each image and retain only the parts corresponding to knot intervals that are completely visible. During the iterations, so as to keep the same cost function, the residual error must be evaluated in the same points in each iteration. Supposing small displacements, we can consider that visible pieces will remain visible throughout the optimization.

The optimization of (8) is done on the 3D control point coordinates, leaving the remaining parameters of the NURBS curve constant. The weights associated with the control points are modified by the projection giving 2D weights varying

with the depth of each control point, according to the formula (4), but they are not subject to the optimization.

In order to avoid over-parameterization for stability reasons, the first optimization is carried out on a limited number of control points. Their number is then increased by iterative insertion, so that the estimated 3D curve fits correctly also high curvature regions. As mentioned earlier, the insertion of a control point is done without influence on the curve and a second optimization is thus necessary to estimate the curve. We finally need a criterion to decide when to stop the control point insertion procedure.

Control Point Insertion. Due to the use of NURBS, we have a method to insert control points. What remains is to decide where to place them. Several strategies have been used. Cham and Cipolla consider in [4] the dual problem of knot insertion. They define an error energy reduction potential and propose to place the knot point so as to maximize this potential. The control point is placed using the method described earlier. In our algorithm, since every insertion is followed by an optimization that locally adjusts the control points, we settle for choosing the interval where to place the point. Since the exact location within the interval is not critical, the point is placed at its midpoint. Dierckx suggests in [5] to place the new point at the interval that presents the highest error. This is consistent with an interpretation of the error as the result of a lack of degrees of freedom that inhibits a good description of the curve. If, however, the error derives from other sources, this solution is not always optimal. In our case, a significant error could also indicate the presence of parasite edges or that of a parallel structure close to the target curve. We have therefore chosen a heuristic approach, that consists in considering all the intervals of the NURBS curve, in order to retain the one that allows for the largest global error decrease.

Stopping Criterion. One of the motives for introducing parametric curves was to avoid treating all curve points, as only the control points are modified during the optimization. If the number of control points is close to the number of samples, the benefit is limited. Too many control points could also cause numerical instabilities, due to an over-parameterization of the curve on the one hand and the size of the non-linear minimization problem on the other hand. It is thus necessary to define a criterion that decides when to stop the control point insertion.

A strategy that aims to avoid the over-parameterization is the use of statistical methods inspired by the information theory. Based in a Maximum Likelihood environment, these methods combine a term equivalent, in the case of a normal distributed errors, to the sum of squares of the residual errors with a term penalizing the model complexity. Given two estimated models, in our case differentiated by their number of control points, the one with the lowest criterion will be retained. The first criterion of this type, called AIC (Akaike Information Criterion), was introduced by Akaike in [1] and is written

$$AIC = 2k + n \ln \frac{RSS}{n}, \quad (9)$$

where k is the number of control points, n is the number of observations and RSS is the sum of the squared residual errors. Another criterion, based on a bayesian formalism, is the BIC (Bayesian Information Criterion) presented by Schwarz [13]. It stresses the number of data points n , so as to ensure an asymptotic consistency and is written

$$BIC = 2k \ln n + n \ln \frac{RSS}{n}. \quad (10)$$

Another family of methods uses the MDL [12] formulation, which consists in associating a cost with the quantity of information necessary to describe the curve. Different criteria follow, depending on the formulation of the estimation problem. In the iterative control point insertion procedure of Cham et Cipolla [4], the stopping criterion is defined by means of MDL. The criterion depends, on the one hand on the number of control points and on the residual errors, on the other hand on the number of samples and on the covariance.

Yet another way of choosing an appropriate model complexity is the classical method of cross-validation. The models are evaluated based on their capacity to describe the data. A subset of the data is used to define a fixed complexity model, while the rest serve to validate it. The process is repeated and a model is retained if its performance is considered good enough.

We have chosen to use the BIC for this first version of our algorithm. A more thorough study of the influence of the stopping criterion in our setting will be performed at a later stage.

4.3 Algorithm

The algorithm we implemented has two parts. The optimization of a curve using a fixed complexity model is embedded in an iterative structure that aims to increase the number of control points. The non-linear optimization of the 3D curve is performed by the Levenberg-Marquardt algorithm, using a cost function based on a search for contour points in the images.

5 Experimental Evaluation

5.1 Virtual Images

In order to validate our algorithms for image data extraction and for curve reconstruction, we have performed a number of tests on virtual images. The virtual setting also allows us to simulate deformations of the target object.

We construct a simplified model of an object, based on a single target curve. We then apply our 3D reconstruction algorithm, starting at a modified “model curve”, on a set of virtual views, see Fig. 1. The image size is 1284×1002 pixels. The starting curve has 10 control points, to which 11 new points are added. The sampling used for the computations is of 200 points. To fix the scale, note that at the mean distance from the object curve, one pixel corresponds roughly to 0.22 mm. The evaluation of the results is done by measuring the distance from

Fig. 1. *Left:* The virtual images used for the reconstruction of the central curve. *Right:* The distances from the sampled points from the reconstructed 3D curve to the model curve. The reconstruction is based on 20 virtual images. The cloud of sample points from the estimated curve is shown together with the target curve. The starting curve is shown in black. The differences are represented by lines with length proportional to the distance between the curve and the target, using a scale factor of 30.

a set of sampled points from the estimated curve to the target model curve. The distances from the target curve are shown in Fig. 1. We obtain the following results:

Mean error	0.0621 mm
Median error	0.0421 mm
Standard deviation	0.0528 mm

We note that the error corresponds to less than a pixel in the images, which indicates a sub-pixel image precision.

5.2 Real Images

We also consider a set of real images, see Fig. 2, with the same target curve, using the same starting “model curve” as in the virtual case. We now need to face the problem of noisy image data, multiple parallel structures and imprecision in the localization and the calibration of the views. The image size is 1392×1040 pixels. The starting curve has 10 control points, to which 11 new points are added. The sampling used for the computations is of 200 points. At the mean distance from the object curve, one pixel corresponds roughly to 0.28 mm. The distances from the target curve are shown in Fig. 2. We obtain the following results:

Mean error	0.157 mm
Median error	0.124 mm
Standard deviation	0.123 mm

Fig. 2. *Left:* Four of the 18 real images used for the reconstruction of the curve describing the central hole. *Right:* The distances from the sampled points from the reconstructed 3D curve to the model curve. The reconstruction is based on 18 real images. The differences are represented by lines with length proportional to the distance between the curve and the target, using a scale factor of 20.

Fig. 3. Problems related to specularities and to the search for candidate points. Starting at the projection of the initial curve (in blue), some candidate points (in magenta) belong to a parasite edge.

Even if the errors are higher than in the case of virtual images, we note that they still correspond to less than a pixel in the images. The difference is partly explained by the noise and the parallel structures perturbing the edge tracking algorithm. An example of candidate points located on a parallel image contour, due to specularities, is given in Fig. 3.

6 Conclusions

We have presented an adaptive 3D reconstruction method using parametric curves, limiting the degrees of freedom of the problem. An algorithm for 3D reconstruction of curves using a fixed complexity model is embedded in an iterative framework, allowing an enhanced approximation by control point insertion. An experimental evaluation of the method, using virtual as well as real images, has let us validate its performance in some simple, nevertheless realistic, cases with specular objects subject to occlusions and noise.

Future work will be devoted to the integration of knowledge of the CAD model in the image based edge tracking. Considering the expected neighborhood of a sample point, the problem of parasite contours should be controlled and has limited impact on the obtained precision.

We also plan to do a deeper study around the stopping criterion used in the control point insertion process, using cross-validation.

References

1. Akaike, H.: A new look at the statistical model identification. *IEEE Transactions on Automated Control* 19(6), 716–723 (1974)
2. Brigger, P., Hoeg, J., Unser, M.: B-spline snakes: A flexible tool for parametric contour detection. *IEEE Trans. on Image Processing* 9(9), 1484–1496 (2000)
3. Cham, T.-J., Cipolla, R.: Stereo coupled active contours. In: *Conference on Computer Vision and Pattern Recognition*, pp. 1094–1099. IEEE Computer Society, Los Alamitos (1997)
4. Cham, T.-J., Cipolla, R.: Automated B-spline curve representation incorporating MDL and error-minimizing control point insertion strategies. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 21(1), 49–53 (1999)
5. Dierckx, P.: *Curve and Surface Fitting with Splines*. Oxford University Press, Inc, New York, USA (1993)
6. Drummond, T., Cipolla, R.: Real-time visual tracking of complex structures. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 7, 932–946 (2002)
7. Hansen, M.H., Yu, B.: Model selection and the principle of minimum description length. *Journal of the American Statistical Association* 96(454), 746–774 (2001)
8. Kahl, F., August, J.: Multiview reconstruction of space curves. In: *9th International Conference on Computer Vision*, vol. 2, pp. 1017–1024 (2003)
9. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active contour models. *International Journal of Computer Vision* 4(1), 321–331 (1987)
10. Meegama, R.G.N., Rajapakse, J.C.: NURBS snakes. *Image and Vision Computing* 21, 551–562 (2003)
11. Piegl, L., Tiller, W.: *The NURBS book*. In: *Monographs in visual communication*, 2nd edn. Springer, Heidelberg (1997)
12. Rissanen, J.: Modeling by shortest data description. *Automatica* 14, 465–471 (1978)
13. Schwarz, G.: Estimating the dimension of a model. *Ann. of Stat.* 6, 461–464 (1978)
14. Siddiqui, M., Sclaroff, S.: Surface reconstruction from multiple views using rational B-splines and knot insertion. In: *First International Symposium on 3D Data Processing Visualization and Transmission*, pp. 372–378 (2002)
15. Xiao, Y.J., Li, Y.F.: Stereo vision based on perspective invariance of NURBS curves. In: *IEEE International Conference on Mechatronics and Machine Vision in Practice*, vol. 2, pp. 51–56 (2001)
16. Yang, H., Wang, W., Sun, J.: Control point adjustment for B-spline curve approximation. *Computer-Aided Design* 36, 639–652 (2004)