

# Contour Segments from Spline Interpolation

Niklas Ludtke and Richard C. Wilson

<sup>1</sup> Beckman Institute, University of Illinois, USA

<sup>2</sup> Dept. of Computer Science University of York, UK

**Abstract.** An important function of perceptual grouping is the restoration of contours. Edge maps produced by low level edge detectors are invariably noisy and inconsistent. It is the aim of perceptual grouping to refine these edge segments by imposing consistency based on considerations about real object outlines. In this paper we describe a method for grouping edge segments into perceptually salient contours using splines. The two important ingredients of our method are firstly the use of probability distributions for possible orientation structure in the image, and secondly the use of Kellman-Shipley relatability to find perceptually meaningful structure. The spline parameters are adjusted to optimise their probabilities in terms of image structure and bending. Consistent structure is then identified using both perceptual criteria and similarity to contour structure in the image.

## 1 Introduction

An important function of perceptual grouping is the restoration of contours. Edge maps produced by low level edge detectors are invariably noisy and inconsistent. It is the aim of perceptual grouping to refine these edge contours by imposing consistency based on considerations about real object outlines. To overcome local distortions in machine vision, numerous authors (e.g. Sha'ashua and Ullman[1]; Sarkar and Boyer[2]; Elder and Zucker[3]; Guy and Medioni[4]) have therefore proposed incorporating contextual relations among local features by combining responses of neighbouring feature detectors into a globalised and consequently more robust processing.

Sha'ashua and Ullman[1] defined a measure of perceptual saliency of a curve, based on geometric properties. The saliency measure increases monotonically with the length of the evaluated curve and decreases with its total squared curvature. Additionally, the degree of fragmentation, expressed in terms of the number of gaps and total gap length, is penalised. A relaxation procedure is then performed to maximise the saliency measure.

Guy and Medioni[4] devised an algorithm for contour grouping, based on the Gestalt principles of co-curvilinearity and proximity. A convolution is performed on the edge map using a special mask called *extension field*, a vector field encoding the likelihood and orientation of possible continuations from an edge segment at its centre to all other points in the image. The direction of the extension field at a point  $(x, y)$  equals the tangent angle of the most likely curve

connecting  $(x, y)$  with the edge segment at the centre of the extension field. The magnitude of the vector field in  $(x, y)$  is the likelihood of the existence of the connecting curve.

Elder and Zucker[3] address the problem of computing closed bounding contours. A multi-scale edge detection algorithm yields information about edge position and tangent orientation from which a sparsely connected tangent graph is constructed. Each node is assigned with the tangent information and, according to a Bayesian model of tangent linking, each arc is labelled with the likelihood that the corresponding tangent pair forms a contiguous component of the same contour. Each node is connected to only a small number of neighbours (usually six), according to the most likely pairings. The goal of closure grouping is then to find the *maximum likelihood cycles* for every tangent in the graph. Thus the grouping task is reduced to a shortest path problem.

Such perceptual organisation is very much in the spirit of Shipley and Kellman’s psychophysically motivated theory of visual interpolation [5]. Although their criterion of edge *reliability* provides a useful test for co-curvilinearity, the geometrical configurations of detected edge segments are in practice often not as precise as required, due to noise. Instead, edge segments would often be erroneously dismissed as “unreliable”, particularly on straight contours. Thus, orientation estimates will often have to be revised, in order to yield “reliable” and more accurate tangent configurations. The revision will be based on mutual consistency, as well as on the quality of agreement between the resultant curve segment and the Gabor transform of the given image.

## 2 Distributions of Orientations and Spline Interpolation

We commence with a description of the orientation structure of the image in terms of a mixture of von-Mises distributions[6, 7].

$$p(\theta) = \sum_i \frac{P_i}{2\pi I_0(\kappa_i)} \exp[\kappa_i \cos(\theta - \theta_i)] \quad (1)$$

This mixture model of von-Mises distributions represents multiple local orientations  $\theta$  and their certainties through the widths  $\kappa_i$ . Typically there will be one or two orientation components, i.e.  $i = \{0, 1\}$ . We therefore encode both multiple directions and uncertainty about each of those directions.

Within this framework, mutual consistency of contours can be expressed in terms of a spline likelihood, comprising the joint density of the orientation pdfs and an additional bending constraint. We have information about the positions and possible tangent directions of edge structure in the image, and therefore we use splines of the quadratic Hermite-type. Their parameters are fully determined by the positions and tangent orientations of the end points. As a consequence of the probabilistic model of orientation, tangent angles are governed by probability densities, and the uncertainty of tangent orientations is transformed into the system’s uncertainty regarding the connecting spline. An important feature of the perceptual grouping framework presented in this paper is that the degree

of uncertainty in the orientation determines the “inertia” of a local tangent estimate, i.e., how easily an initial orientation measurement (given by a mode in the corresponding pdf) can be modified during consistency optimisation. In this role of certainty lies a conceptual difference to other grouping schemes. Usually, the coarseness of the initial local orientation measurements is expressed in terms of the likelihood of potential continuations at the grouping level, for example, characterised by a “support function” in relaxation labelling [8, 9] or by orientation “votes” propagated through an “extension field” [4]. The *initial* certainty of the local measurement, however, is not modelled.

### 3 Quadratic Splines

After low-level processing[6, 7] the positions of the control points, the corresponding distributions of tangent orientations are known, and points with multiple orientations are identified as such. Therefore, *piecewise quadratic spline interpolation* provides a very straightforward means of connecting such control points.

In the grouping framework presented, the constraint of  $C^2$  and  $C^1$  continuity at the control points will not be imposed. Instead, left and right limits of tangent orientation are introduced, whereby smooth and polygonal curves can be represented equally well [10]. As a result, the algorithm is capable of representing tangent discontinuities suggested by feature associations, in addition to the locally detected points with multiple orientations.

Let  $t \in [0, 1]$  be the spline parameter, and let  $\mathbf{s}(t)$  denote a position on the spline, i.e.,  $\mathbf{s}(t) = [x(t), y(t)]^T$ . Then the quadratic spline is defined as:

$$\mathbf{s}(t) = \mathbf{a}t^2 + \mathbf{b}t + \mathbf{c}, \quad \text{with } \mathbf{a}, \mathbf{b}, \mathbf{c} \in \mathbb{R}^2. \quad (2)$$

Here  $\mathbf{a}$ ,  $\mathbf{b}$  and  $\mathbf{c}$  are the vector-valued spline coefficients. These quantities are not geometrically meaningful in this form, so we write the spline in terms of the positions of the endpoints  $\mathbf{r}_1, \mathbf{r}_2$  and the tangent directions at the endpoints;  $\theta_1, \theta_2$ .

$$\mathbf{s}(t) = (-t^2 + 1)\mathbf{r}_1 + t^2\mathbf{r}_2 + (-t^2 + t) \cdot 2d \frac{\sin(\phi - \theta_2)}{\sin(\theta_1 - \theta_2)} \begin{pmatrix} \cos \theta_1 \\ \sin \theta_1 \end{pmatrix} \quad (3)$$

Here  $d$  and  $\phi$  are the length and angle of the vector connecting the endpoints,  $\mathbf{r}_2 - \mathbf{r}_1$ . It is important to note that the final term becomes singular for  $\theta_1 = \theta_2$ . Since the splines do not contain inflexions, equality of  $\theta_1$  and  $\theta_2$  is only possible if  $\theta_1 = \theta_2 = \phi$ . The spline therefore approaches the straight line  $\mathbf{s}(t) = \mathbf{r}_1 + t(\mathbf{r}_2 - \mathbf{r}_1)$ .

### 4 Optimisation of Spline Parameters

The essential point of our method is that the local orientation is represented by a probability distribution, and therefore the spline parameters  $\theta_1$  and  $\theta_2$  are not

fixed quantities, and may be varied in order to obtain more consistent splines. The final splines may be reliable even if the initial configuration  $(\bar{\theta}_1, \bar{\theta}_2)$  is not.

By virtue of (3), each pair  $(p(\theta_1), p(\theta_2))$  of two locally extracted orientation densities implies a density  $p(\mathbf{s}|\theta_1, \theta_2)$  in the spline, thus describing a “bundle” of possible quadratic splines passing through the fixed end points. There is no need to actually compute  $p(\mathbf{s}|\theta_1, \theta_2)$ . Instead, the optimisation is performed with respect to the tangent angles  $\theta_1$  and  $\theta_2$ , and the corresponding optimal  $\mathbf{s}$  is calculated afterwards.

In order to enforce smoothness of contours, it is necessary to impose a shape constraint on the connecting spline bundle that penalises a high degree of bending. The new tangent angles are then found by means of a maximum likelihood estimation procedure, which results in a trade-off between closeness to initial local measurements and smoothness constraint. The final decision about the reliability of a pair of key points is made after this optimisation.

#### 4.1 The Spline Likelihood Function

In general terms, the total likelihood of a pair of tangent angles  $(\theta_1, \theta_2)$  is given by the product of the joint density of that pair, obtained from (1), and a probability density that depends on the degree of bending of the corresponding spline:

$$\mathcal{L}(\theta_1, \theta_2) = p_{\text{pop}}(\theta_1, \theta_2) p_{\text{bend}}(\theta_1, \theta_2). \quad (4)$$

The quantity that describes the bending of the spline is the magnitude of the second derivative with respect to the parameter  $t$ , denoted by  $\ddot{s}(t)$ . For a quadratic spline  $\ddot{s}$  is a constant determined by the boundary conditions. The “bending probability”, which acts as a penalty term in the likelihood function (4), can be defined as a Gaussian distribution in the scale invariant quantity  $\ddot{s}/d$ :

$$p_{\text{bend}}(\theta_1, \theta_2) = p(\ddot{s}(\theta_1, \theta_2)) = \frac{1}{\sqrt{2\pi}\sigma_b} \exp\left(-\frac{\ddot{s}^2(\theta_1, \theta_2)}{2d^2\sigma_b^2}\right) \quad (5)$$

The variance  $\sigma_b^2$  determines how strongly bending is penalised and is to be optimised together with  $\theta_1$  and  $\theta_2$ . In order to find  $\ddot{s}(\theta_1, \theta_2)$ , equation (3) is differentiated twice with respect to  $t$ , yielding the square of its magnitude as:

$$\ddot{s}^2(\theta_1, \theta_2) = 4d^2 - 16d^2 \frac{\sin(\phi - \theta_2) \cos(\phi - \theta_1)}{\sin(\theta_1 - \theta_2)} + 16d^2 \frac{\sin^2(\phi - \theta_2)}{\sin^2(\theta_1 - \theta_2)}$$

The density relating to the image orientation may be expressed as the product of the densities for each endpoint, since the local orientations are treated as independent random variables:

$$\begin{aligned} p_{\text{pop}}(\theta_1, \theta_2 | i, j) &= p(\theta_1 | i) p(\theta_2 | j) \\ &= \frac{1}{4\pi^2 I_0(\kappa_i) I_0(\kappa_j)} e^{\kappa_i \cos(\theta_1 - \bar{\theta}_i) + \kappa_j \cos(\theta_2 - \bar{\theta}_j)} \end{aligned} \quad (6)$$

In most cases both points are on an edge, and consequently  $i = 1$  and  $j = 1$ . Only in corner points or junctions several possible associations  $(i, j)$  have to be considered. Inserting (5) and (6) in (4) the spline log-likelihood becomes:

$$\ln \mathcal{L}(\theta_1, \theta_2; i, j) = -\ln[4\pi^2 I_0(\kappa_i) I_0(\kappa_j)] + \kappa_i \cos(\theta_1 - \bar{\theta}_i) + \kappa_j \cos(\theta_2 - \bar{\theta}_j) - \frac{1}{2} \ln[2\pi] - \ln \sigma_b - \frac{\dot{s}^2(\theta_1, \theta_2)}{2d^2 \sigma_b^2} \quad (7)$$

The log-likelihood may then be solved using a standard optimisation method. For the results presented in this paper, we used a simple gradient ascent approach.

It is important to note that in the log-likelihood function (7) the concentration parameters  $\kappa_i$  and  $\kappa_j$  of the pdfs act as weights of angular modifications during the optimisation process. If a concentration parameter is large, any deviation from the initial orientation  $\bar{\theta}$  will result in a sharp decrease of the likelihood function unless the overall curvature is substantially reduced simultaneously. In other words, the concentration parameters (and thus the certainties, which are monotonic functions thereof) determine the “inertia” of orientation estimates, i.e., their “flexibility to compromise for the sake of mutual consistency”. Herein lies the essential difference to other grouping methods, where measurement of certainty is not an integral part of local feature extraction.


## 4.2 Detection of Control Points

To implement such a spline scheme, we must begin by locating the necessary control points. These control points are located at points of significant edge response in the image. Given a set of key points, the task is then to decide which points can be connected by splines. Here the idea is to eliminate only very unlikely configurations, leaving more plausible arrangements to be decided on after the splines have been reconfigured. For each key point, only a limited number of its nearest neighbours are considered for grouping, reflecting the Gestalt law of *proximity*. A weak relatability criterion is then applied to remove connections which are inconsistent. After optimisation of the spline parameters, the resulting splines are checked to establish whether they correspond to real edge structure in the image. Those that have no support are discarded. Furthermore, Shipley-Kellman relatability is checked with regard to the splines neighbours. Those that are inconsistent are discarded.


## 5 Experiments

Figure 1 shows an image of part of the sculpture of Paolina Borghese by Antonio Canova (1757-1822) and the different steps of contour extraction, from Gabor responses to the tangent elements extracted at key points and the result of spline interpolation between them. Some parts of the hair region contain very narrow features which are more suitable for processing with a line detector. The method has difficulties in correctly relating segments which are parallel and very close to each other.

The same image has been used by Iverson and Zucker (1995) to demonstrate the performance of their “logical/linear operators”. Referring to earlier work by




**Fig. 1.** Stages of contour extraction using a photograph of the sculpture “Paolina” ( $512 \times 512$  pixels, from the archive of the Vision group of Pietro Perona at Caltech).




**Fig. 2.** The effect of tangent optimisation. In (a) the “reliability” criterion of Shipley and Kellman (1991) is applied directly to the tangent orientations given by the modes of the corresponding mixture densities. Figure (b) shows the same process after tangent optimisation.

Koendrink and co-workers (1982), the authors point out the perceptual significance of bifurcations and line terminations in regions, such as the folds around the neck, which provide vital information about three-dimensional structure. They also demonstrate that the Canny detector [13], like any other essentially linear edge operator, is not capable of correctly representing bifurcations and tends to smooth out tangent discontinuities in corner points and T-junctions. Since feature extraction with probabilistic population coding explicitly represents points with multiple orientations and orientation discontinuities, the spline interpolation algorithm can accurately capture most of the essential discontinuities and bifurcations (Fig. 1 (f)).

Figure 2 illustrates the effect of the tangent optimisation algorithm on the spline contour representation. When applied directly to the tangent orientations given by the modes of the corresponding (mixture) densities, the “reliability” criterion of Shipley and Kellman (1991) rejects a number of tangent configurations, and consequently a lot of contour segments are not detected. Also, many



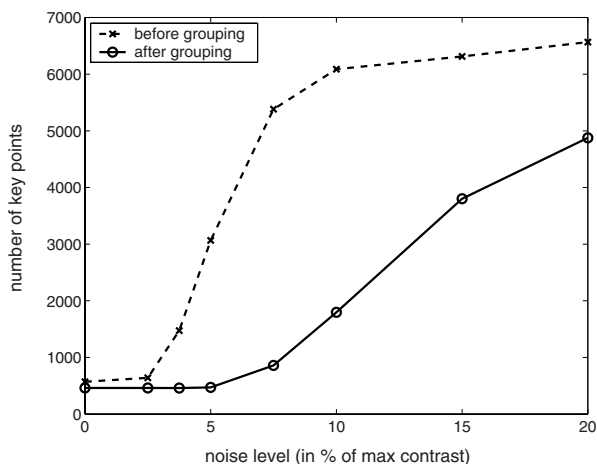
**Fig. 3.** Original image with the tangent map and spline contours.



**Fig. 4.** An image with different amounts of additive Gaussian noise.

splines tend to differ from the actual contours, since inexact tangent angles tend to create curved rather than straight splines. Through optimisation, a significant number of tangent angles can be adjusted, in order to yield “reliable” configurations, many of which prove to be consistent with the intensity gradient in the image. As a result, a more complete and accurate contour representation is obtained.

In another experiment, the performance of the algorithm in the presence of additive Gaussian has been investigated. Figure 3 shows the original image with the extracted spline contours. Figure 4 shows the result of feature detection and subsequent perceptual grouping for moderate noise. For a moderate noise level ( $\sigma_N < 5\%$ , SNR  $< 26$  dB) there are only few false positives in the spline representation, since most erroneous key points form only isolated splines that can easily be identified and removed. Above a noise value of about 10% (SNR = 20 dB), the density of false positive key points reaches a level where spurious splines begin to form erroneous contour segments of considerable length which could only be eliminated by perceptual organisation of higher order. At this stage curvature



**Fig. 5.** The number of key points detected in Figure 4 (a) as a function of the noise level (standard deviation in % of maximum contrast) before (dashed curve), and after perceptual organisation (solid curve).

consistency would be a vital constraint, since the noise-induced contour segments exhibit frequent, sudden changes in the sign of curvature, which rarely occur in natural object boundaries and folds. The erroneously *discarded* key points (false negatives) are small in number but, of course, much more obvious, since they lead to gaps in the contour representation.

The number of detected features as a function of the noise level is an indicator for the efficiency of the feature detection in the presence of noise, since the number of additional key points compared to the case without noise approximately equals the number of false positives. Figure 5 shows a plot of this relation for the image in Figure 4 (a). Though more and more spurious key points appear with increasing noise level, most of them do not fulfill the reliability criterion, and even after angular optimisation a potential spline connection often lacks consistency with the filter responses. Thus many false positives can be identified and rejected.

## References

1. Sha'ashua, A., Ullman, S.: Structural saliency: The detection of globally salient structures using a locally connected network. In: Proceedings of the 2nd ICCV. (1988) 321–327 Also Weizmann Institute of Science Report CS88-18, (October 1988).
2. Sarkar, S., Boyer, K.: Computing Perceptual Organisation in Computer Vision. World Scientific (1994)
3. Elder, J., Zucker, S.: Computing contour closure. In: Proceedings of the European Conference on Computer Vision. Volume Lecture Notes in Computer Science 1064., Springer Verlag (1996) 399–



4. Guy, G., Medioni, G.: Inferring global perceptual contours from local features. *Int. J. Comp. Vis.* **20** (1996) 113–133
5. Kellman, P., Shipley, T.: A theory of visual interpolation in object perception. *Cognitive Psychology* **23** (1991) 141–221
6. Lüdtkke, N., Wilson, R., Hancock, E.: Population codes for orientation estimation. In: *Proceedings of the 15th International Conference on Pattern Recognition.* (2000)
7. Lüdtkke, N., Wilson, R., Hancock, E.: Probabilistic population coding of multiple edge orientation. In: *Proceedings of the International Conference on Image Processing, IEEE* (2002) to appear.
8. Parent, P., Zucker, S.: Trace inference, curvature consistency, and curve detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **11** (1989) 823–839
9. Hancock, E., Kittler, J.: Discrete relaxation. *Pattern Recognition* **23** (1990) 711–733
10. Gavrilu, D.: Hermite deformable contours. Technical Report CS-TR-3610, Center for Automation Research, University of Maryland (1996) also ICPR'96.
11. Iverson, L., Zucker, S.: Logical/linear operators for image curves. *IEEE Transactions on Pattern Recognition and Machine Intelligence* **17** (1995) 982–996
12. Koendrink, J., van Doorn, A.: The shape of smooth objects and the way contours end. *Perception* **11** (1982) 129–137
13. Canny, J.: A computational approach to edge detection. *IEEE Transact. on Patt. Rec. and Machine Intell.* **8** (1986) 679–700