

A Statistical Process for Surface Tracking

Nigel G. Sharp and Edwin R. Hancock

Department of Computer Science, University of York, YO1 5DD, UK

Abstract. This paper describes a novel approach to surface tracking in volumetric image stacks. It draws on a statistical model of the uncertainties inherent in the characterisation of feature contours to compute an evidential field for putative inter-frame displacements. This field is computed using Gaussian density kernels which are parameterised in terms of the variance-covariance matrices for contour displacement. The underlying variance model accommodates the effects of raw image noise on the estimated surface normals. The evidential field effectively couples contour displacements to the intensity features on successive frames through a statistical process of contour tracking. Hard contours are extracted using a dictionary-based relaxation process. The method is evaluated on both MRI data and simulated data.

1 Introduction

Multi-frame feature tracking is an essential ingredient in both motion analysis and volumetric image reconstruction. The basic task involves the robust identification and location of feature contours that exhibit inter-frame deformations caused either by object motion or by significant surface curvature. In order to be effective, tracking must be capable of distinguishing genuine variations in both the position, shape and intensity when the available image information is noisy or uncertain. Existing techniques described in the literature can be divided into those that attempt to capture contour deformation at either the global or local levels. Global methods [1, 8] aim to locate acceptable parametric contour deformations that satisfy constraints on consistent inter-frame movement [1] and closeness to the raw intensity data [8]. However, their main limitation stems from the need for prior knowledge governing the topology of the target contours[1]. If this knowledge is not to hand, then local tracking methods offer an attractive alternative which also effectively draw on information concerning intensity structure [5, 7] and contour consistency [5]. However, while the global methods accommodate this information in a unified optimisation framework, local feature tracking methods conventionally draw upon it sequentially in a multistage approach.

The localised multistage philosophy is nicely exemplified in the volumetric imaging domain. Here MRI slice analysis is invariably approached as a three stage segmentation strategy [3, 11]. Intensity features are first enhanced by volumetric filtering, prior to linking and surface fitting. Each of these stages requires an intermediate representation of the data which effectively discards important information. By contrast relaxation processes offer a unified framework in

which local information may be combined to produce a consistent description of underlying contour structure [5, 6]. As a concrete example, we have recently developed a novel relaxation framework [4] which has proved extremely robust when applied to 2D feature detection problems [5, 6, 7]. Much of its robustness stems from the use of contextual information, in particular a dictionary representation of 2D contour structures [5, 4], in the refinement of intensity-features whose uncertainties are characterised by probability density functions. These probability densities directly model the effects of applying image operators to the raw image data. In consequence, their parameters may be directly estimated from the raw-image statistics, endowing the relaxation approach with adaptive capabilities [6].

The aim of this paper is to extend these ideas by developing a relaxation process for multi-frame curve tracking [11, 9]. The motivation for this study stems from the need to directly exploit local constraints on inter-frame contour deformation, avoiding the need for global surface models of predefined topology. In other words, since we wish to avoid the well documented difficulties associated with the inability of global contour models to adapt to changes of topology, we appeal to a local curve evolution process [1, 2, 10, 12]. To some extent our model is analogous with the reaction-diffusion process for curve-evolution [2, 12, 10]. According to this picture the dynamics of curve evolution may be classified according to the equations of motion in the normal-direction. Under the reaction limit curves evolve with uniform speed in a manner which has been likened to morphological transformation under the action of a distance map [2]. In the diffusion case the evolution process is analogous to wavefront propagation with curvature dependant speed [12].

Motivation for our study comes from the need to exploit the rich sources of evidence available for controlling consistent curve evolution. Taking the tracking of contours in volumetric image stacks as an example, this exists not only in the form of statistical information derived from the raw intensity image but also in terms of constraints on local curve evolution [13, 14]. In this paper we demonstrate how these various sources of information can be combined or fused together to compute a unified evidential index for consistent inter-frame contour evolution. This index is statistical in character and can be thought of as a probabilistic counterpart of the morphological distance map used in the reaction-diffusion analysis of curves [2]. Central to this fusional process is the availability of an evidence combining framework [5]. The approach adopted here is a novel relaxation process that is tailored to multi-frame feature tracking [14] and allows constraints to be imposed on consistent contour development. As pointed by Zucker *et. al.* [15] relaxation processes may be thought of as a local form of snake dynamics. The aim of our work is therefore to provide a statistical process for local contour tracking that can couple to image features via a distance map.

The outline of this paper is as follows. Section 2 reviews our multi-frame relaxation process for contour tracking. Section 3 develops our surface tracking model and identifies its relationship with alternative models of contour evolution. In section 4 we evaluate the method on volumetric image data. Finally, Section 5 provides some conclusions.

2 Multi-frame Relaxation

Our adopted framework for surface tracking is the multi-frame relaxation technique described in [14]. Before we proceed to detail the modelling of contour evolution, we describe some of the formal elements of the relaxation framework. We commence by specifying the form and nature of the information available to the tracking process from the different frames in the sequence. Suppose that the image entities forming the frame with index L are denoted by the set F^L ; this set indexes the pixel co-ordinates in the frame of interest. We will be interested in labelling the pixels in the frame F^L by combining evidence in the form of within-frame measurements together with labelled feature contours from the previous frame in the sequence, i.e. F^{L-1} . We let \underline{x}_j^L denote the measurement-vector pertaining to the pixel $j \in F^L$. Typically, this vector represents the output of a series of feature detection filters [7].

2.1 Bayesian Ingredients

The basic philosophy in developing the multi-frame relaxation formula is to combine evidence from the 3×3 neighbourhood I_j^L of pixel j in the light of constraints on consistent contour evolution. When the relaxation operations have converged, each pixel in the frame F^L has a hard interpretation according to the set of available feature labels, Ω . We denote this labelling by the set $\Upsilon^L = \{\theta_j^L \in \Omega, \forall j \in F^L\}$. This hard labelling may be utilised as an additional source of contextual information to track features in the subsequent frames in the sequence. Additional information of this sort can only effectively enhance the quality of interpretation provided that a suitable model of inter-frame feature persistence is available. It is this model requirement that places the greatest demands on our multi-frame relaxation method and is the main topic of this paper.

The development of the multi-frame relaxation formula [14] draws on the above ingredients using the following non-linear probability update formula

$$P^{(n+1)}(\theta_j^L) = \frac{P^{(n)}(\theta_j^L)Q^{(n)}(\theta_j^L)}{\sum_{\theta_j^L \in \Omega} P^{(n)}(\theta_j^L)Q^{(n)}(\theta_j^L)} \quad (1)$$

where $P^{(n)}(\theta_j^L)$ is the label probability for pixel j in frame L at iteration n of the relaxation scheme. Evidence for the label assignment θ_j^L is accumulated over the contextual neighbourhood I_j^L in frame F^L by the support function $Q^{(n)}(\theta_j^L)$. The critical Bayesian ingredient of the multi-frame support function are the inter-frame priors $P(\Upsilon^L | \Upsilon^{L-1})$ which represent constraints on the evolution of consistent contour structure. These constraints are modelled by compiling those contours on frame F^L that are consistent with their hard labelled counterparts in frame F^{L-1} , i.e. Υ^{L-1} in a dictionary $\Theta(\Upsilon^{L-1})$. If λ_k^h is the label on the object indexed k , corresponding to the h th entry in the dictionary $\Theta(\Upsilon^{L-1})$, then we can introduce the following shorthand notation for the dictionary items




Fig. 1. Contour dictionary

$\Lambda^h = \{\lambda_k^h, \forall k \in I_j^L\}$. In practice the dictionary is constructed so as to discourage contour fragmentation or radical changes in feature orientation.

With the dictionary ingredient, the support function for multi-frame feature tracking is

$$Q^{(n)}(\theta_j^L) = \frac{1}{P(\theta_j^L | \Upsilon^{L-1})} \times \sum_{\Lambda^h \in \Theta(\Upsilon^{L-1})} \left\{ \prod_{k \in I_j^L - j} \frac{P^{(n)}(\lambda_k^h)}{P(\lambda_k^h | \Upsilon^{L-1})} \right\} P(\Lambda^h | \Upsilon^{L-1}) \quad (2)$$

This formula is initialised by setting $P^{(0)}(\theta_j^L) = P(\theta_j^L | \mathbf{x}_j^L, \Upsilon^{L-1})$. From the Bayesian standpoint, this means that the initial probabilities depend both on the measurements pertaining to the current frame and upon hard labelled contour information derived from the previous frame.

2.2 Dictionary Model

For the 2D feature detection application the dictionary consists of all the connected contours which can be constructed on a 3x3 neighbourhood using the label-set $\Omega = \{\leftrightarrow, \updownarrow, \phi\}$. There are about 100 such configurations; some typical examples are shown in Figure 1. All the consistent contours are a single pixel wide. The dictionary contains symbolic representations of important feature points such as corners and T-junctions.

At this point it is worth contrasting the philosophy underlying our relaxation approach with that adopted in deformable template models[1]. Rather than addressing the inter-frame evolution of contours in terms of global parameter

deformations, we attempt to identify consistent structure that satisfies local constraints. The role of the conditional priors is therefore to model the development of acceptable persistent contour structure which meets certain inter-frame consistency requirements. Our modelling of these requirements is based on the dictionary items satisfying constraints on acceptable inter-frame variation in contour orientation and curvature. We therefore partition the dictionary items on frame L into sets deemed acceptable according to the consistent labellings obtained on frame $L - 1$; we denote this acceptable set by $S(\Upsilon^{L-1})$. In order to compute the inter-frame conditional priors, we apportion the available probability mass uniformly among the dictionary items belonging to the partition $S(\Upsilon^{L-1})$ i.e.,

$$P(\Upsilon^L | \Upsilon^{L-1}) = \begin{cases} \frac{1}{|S(\Upsilon^{L-1})|} & \text{if } \Upsilon^L \in S(\Upsilon^{L-1}) \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

With the joint conditional priors to hand, the single object conditional priors $P(\theta_j^L | \Upsilon^{L-1})$ may be computed from the relative frequencies of the various object labels in the partitioned dictionary. In doing this we implicitly assume that the single object priors are independent of the bulk of contour information contained in the slice $L - 1$, i.e. $P(\theta_j^L | \Upsilon^{L-1}) = P(\theta_j^{L-1})$. This is a physically plausible assumption; it means that *a priori* we have no prejudices about the location of feature and non-feature points, it is only when evidence is to hand that we modify our expectations.

The multi-frame relaxation formula provides a powerful framework for feature tracking which can simultaneously accommodate inter-frame contour deformations, statistical intensity variations, and, the local differential structure of intensity surfaces. Conventionally the tracking process is approached by either the construction of deformable contour models [2, 10] or via predictive statistical models such as Kalman filters [1]. In our approach the statistical origins of feature variability are captured at three levels. Intrinsic intensity variations are described by a Gaussian noise model. Contour shape evolution due to deformations and movements of the underlying surface are captured by a statistical model of the contour evolution process. Constraints on local contour structure are used to impose consistency on the detected surface traces.

3 Statistical Evolution of Contours

In this paper we incorporate statistical information concerning local surface structure into the underlying tracking model. Ultimately, we would like to draw on all the information residing in the local Darboux frame [13], i.e. the surface normal together with the magnitudes and directions of principal curvature. However, since curvature computation is severely limited by noise, for the time being we will concentrate on modelling the statistical uncertainties present when surface normals are estimated using the set of filter responses.

The underlying philosophy is to compute an evidential field for plausible contour displacements. The computation of this field is effected using anisotropic Gaussian density kernels whose parameters are the contour displacement variances. We adopt a variance model in which accommodates the effect of raw image

noise on the computed surface normals. Specifically, we associate an evidential field with plausible contour deformations. This field is subsequently refinement by relaxation operations.

3.1 Surface Normals

Suppose that the volumetric intensity information at the point with co-ordinates x, y, z is denoted by $I(x, y, z)$. We are concerned with modelling the evolution of curves in the x-y plane as the stack is descended in the z-direction. The kernels for feature detection in the directions of the three co-ordinate axes are $K_x(x, y, z)$, $K_y(x, y, z)$ and $K_z(x, y, z)$. These three kernels are convolved with the raw volumetric intensity data to compute a vector of filter responses $\vec{N} = (F_x(x, y, z), F_y(x, y, z), F_z(x, y, z))^T$ which is used to approximate the direction of the surface normal. Under quite non-restrictive assumptions, it can be shown that if the raw image data is subject to additive Gaussian noise of zero mean and variance σ^2 , then the noise distribution for the surface normal is a multivariate Gaussian of zero mean. The elements of the variance-covariance matrix Σ_N are related to the autocorrelation f_u and cross correlation f_c of the filter kernels used to characterise the surface intensity features. Specifically, when the three kernels are rotated versions of one-another,

$$\Sigma_N = \sigma^2 \begin{pmatrix} f_u & f_c & f_c \\ f_c & f_u & f_c \\ f_c & f_c & f_u \end{pmatrix} \quad (4)$$

In our experimental evaluation of the tracking process, we will be interested in detecting surfaces associated with volumetric edges. Here we use both overlapping and non-overlapping derivative of Gaussian filters. Since the kernels are antisymmetric their cross-correlations are zero i.e. $f_c = 0$.

In order to realise tracking as a statistical process we will be interested in computing the distribution of contour movements between adjacent slices of the volumetric image stack. Figure 2a illustrates how we calculate the expected contour displacement \vec{D} . Here, \vec{N} is the surface normal approximated by the filter responses and \vec{T} is a surface tangent vector leading from the current contour point to the corresponding contour point on the next slice. The vector \vec{T} is perpendicular to \vec{N} , and is chosen such that its projection onto the x-y plane \vec{D} is directed along \vec{N} 's projection onto the x-y plane. These constraints lead to the following equations for calculating the inter-slice contour displacement \vec{D} , given the filter responses F_x, F_y, F_z and the inter-slice spacing t :

$$D_x(x, y, z) = -\frac{tF_x(x, y, z)F_z(x, y, z)}{F_x(x, y, z)^2 + F_y(x, y, z)^2} \quad (5)$$

$$D_y(x, y, z) = -\frac{tF_y(x, y, z)F_z(x, y, z)}{F_x(x, y, z)^2 + F_y(x, y, z)^2} \quad (6)$$

The displacements are directed along the component of the surface normal in the x-y plane and are proportional to the slope of the tangent plane of the local intensity surface.




Fig. 2. a) Relationship of surface normal \vec{N} , tangent vector \vec{T} , and inter-slice contour displacement vector \vec{D} ; b) Ellipsoidal contour distribution.

Our basic aim in this paper is to compute an evidential field for contour displacements by drawing on a set of probability densities that characterise the various pieces of information derived from the image stack. For representational convenience we therefore amalgamate the vector of filter responses with the vector of contour displacements $\vec{\Delta}_j^L = (D_x(x, y, z), D_y(x, y, z))^T$ to compute a combined feature-vector $\mathbf{x}_j^L = (\vec{N}, \vec{\Delta}_j^L)^T$ for the pixel indexed j . With these ingredients we can compute the conditional density function for pixels belonging to the background class ϕ which do not constitute consistent contours resulting from meaningful intensity surfaces. This is effectively the noise-density for structureless image regions and is given by

$$p(\mathbf{x}_j^L, \Upsilon^{L-1} | \theta_j^L = \phi) = \frac{1}{(2\pi)^{\frac{3}{2}}} \frac{1}{\sqrt{|\Sigma_N|}} \exp \left[-\frac{1}{2} \vec{N}^T \Sigma_N^{-1} \vec{N} \right] \quad (7)$$

The conditional density function for structured image regions is more elusive. In the next section we will describe an ansatz for computing this density for contours which evolve as the image stack is descended. As a prerequisite we require a covariance matrix to represent the underlying uncertainties in contour displacement.

3.2 Contour Displacement Uncertainties

We are interested in providing the resulting field of inter-frame contour displacements with a probabilistic description. For this reason, we must direct our attention to understanding how the statistical uncertainties in the components of surface normals propagate through into the computation of contour displacement. The propagation of errors is controlled by the Jacobian of the transformation between normal-components and contour displacements given in equations (2) and (3), *i.e.*.

$$J = \frac{t}{(F_x^2 + F_y^2)^2} \begin{pmatrix} F_z(F_x^2 - F_y^2) & 2F_x F_y F_z \\ 2F_x F_y F_z & F_z(F_y^2 - F_x^2) \\ -F_x(F_x^2 + F_y^2) & -F_y(F_x^2 + F_y^2) \end{pmatrix} \quad (8)$$

With the Jacobian to hand, the transformed variance-covariance matrix for the contour displacements is equal to $\Sigma_D = J^T \Sigma_N J$. After matrix manipulation and

algebra the contour displacement covariance matrix is equal to

$$\Sigma_D = \frac{t^2 f_u \sigma^2}{(F_x^2 + F_y^2)^2} \begin{pmatrix} F_x^2 + F_z^2 & F_x F_y \\ F_x F_y & F_y^2 + F_z^2 \end{pmatrix} \quad (9)$$

The above covariance matrix captures statistical uncertainties in the displacement field due to the effects of raw image noise on the computation of the tangent slope.

3.3 Contour Mixture Density

Key to meeting our objective of tracking evolving contours is a means generating an evidential field from the hard labelled contours in preceding frames of the volumetric image stack. This softening process effectively disperses the detected contour points about the computed direction of the surface normal. Our basic assumption in formulating the softening process is that the labelled feature pixels in the preceding frame of a sequence represent mutually exclusive seeds from which putative feature points on the current frame may have developed. We may therefore compute the conditional measurement density for the tentative feature point at the pixel indexed j in frame F^L by summing over the contributing densities from the hard-labelled feature points in frame F^{L-1} . The resulting mixture density is given by

$$p(\underline{x}_j^L, \Upsilon^{L-1} | \theta_j^L = \epsilon) = \sum_{k \in F^{L-1}} p(\underline{x}_j^L, \theta_k^{L-1} = \epsilon | \theta_j^L = \epsilon) P^{(\infty)}(\theta_k^{L-1} = \epsilon) \quad (10)$$

where ϵ represents a combined horizontal and vertical feature label. The above expression for the conditional density effectively blurs the hard-labelled contours from frame F^{L-1} to produce an evidential field for potential contours in frame F^L . The hard label information is represented by $P^{(\infty)}(\theta_k^L = \epsilon)$ which indicates whether or not the pixel k is assigned to the feature class in frame $L-1$. Evidential information is represented by the individual densities $p(\underline{x}_j^L, \theta_j^{L-1} = \epsilon | \theta_k^L = \epsilon)$ for the features contained within frame F^L to have originated from contours in frame F^{L-1} . It is these densities that form the critical and novel ingredient of our contour tracking method. They model the evolution of tentative feature points on the frame F^L from the labelled feature points on the frame F^{L-1} . Their role is to capture the spatial variations in contour position on different frames. In so doing, they account for inter-frame deformation of 2D contours originating from the surface geometry of 3D structures.

We would like to utilise the computed uncertainties in the displacement components to construct probability density functions to model the contour evolution process. Since we have little information concerning the prior distribution of surface structure to hand, the construction of accurate distributions is not feasible. For convenience, we therefore associate Gaussian densities with the distribution of contour displacements. If \vec{q}_j^L is a vector of x-y co-ordinates for the pixel indexed j in frame L , i.e. $\vec{q}_j^L = (x_j, y_j)^T$, then the required density must

effectively compare the putative contour displacement $\vec{\delta}_{j,k}^L = \vec{q}_j^L - \vec{q}_k^{L-1}$ with the predicted displacement $\vec{\Delta}_j^L$, i.e.

$$p(\underline{x}_j^L, \theta_k^{L-1} = \epsilon | \theta_j^L = \epsilon) = \frac{1}{2\pi} \frac{1}{\sqrt{|\Sigma_D|}} \exp \left[-\frac{1}{2} (\vec{\delta}_{j,k}^L - \vec{\Delta}_j^L)^T \Sigma_D^{-1} (\vec{\delta}_{j,k}^L - \vec{\Delta}_j^L) \right] \quad (11)$$

In this way the detected contour points in frame $L - 1$ are distributed in an ellipsoidal fashion in frame L . This ellipsoid is centred according to the predicted displacement and aligned along the projected normal direction in the x-y plane as illustrated in Figure 2b. The net effect of this statistical tracking process is to generate a probabilistic counterpart of the weighted distance-map. The density appearing in equation (7) may be converted into an *a posteriori* probability by a simple application of the Bayes formula

$$P(\theta_j^L = \epsilon | \underline{x}_j^L, \Upsilon^{L-1}) = \frac{p(\underline{x}_j^L, \Upsilon^{L-1} | \theta_j^L = \epsilon) P(\theta_j^L = \epsilon)}{p(\underline{x}_j^L, \Upsilon^{L-1} | \theta_j^L = \epsilon) P(\theta_j^L = \epsilon) + p(\underline{x}_j^L, \Upsilon^{L-1} | \theta_j^L = \phi) P(\theta_j^L = \phi)} \quad (12)$$

The resulting probability distribution effectively provides a unified way of coupling feature characteristics derived from the filtering of frame L with the displacement field derived from hard-labeled contours on frame $L - 1$; it can be regarded as a providing a combined index for contour evolution between successive frames. Hard labelled contours may be identified by applying the apparatus of Bayes decision theory to the probabilities defined in equation (9). In our experimental studies we will draw on the multi-frame relaxation scheme described in [14].

4 Experiments

We have evaluated our tracking method on both MRI data and synthetically generated volumetric images. The synthetic data has been created with a view to assessing the capabilities of our method in the tracking of contours associated with highly inclined and curved surfaces. It also serves to illustrate the advantages to be gained from statistical model of local surface structure over our earlier and more naive tracking process [14] which draws only on the proximity of adjacent contours. Firstly, Figure 3a shows a sequence of MRI scans of a brain tumor, sampled vertically from a volumetric image stack. Figure 3b shows the initial evidential field computed by statistical surface tracking, i.e. the value of $P^{(0)}(\theta_j^L = \epsilon)$, prior to the application relaxation operations. The sequence of images in Figure 3c shows the result of applying the relaxation operator to the computed initial probabilities. It is interesting to note that the initial probabilities are well localised; the statistical surface tracking process is well collimated and does not, for instance, disperse or scatter the available evidence in an incoherent way. The hard labelled contours maintain good contact with the underlying edge features and curvature details .

To re-enforce some of these points, Figure 4a shows a series of images from a synthetic volume containing two quarter-segments of a sphere. The raw sequence




Fig. 3. Top (a) MRI sequence; Middle (b) Statistical contour field; Bottom (c) Labelled contours.

contains additive Gaussian noise such that the signal-to-noise ratio for the genuine edge features is 5:3. The initial evidential fields are shown in Figure 4b while the hard labelled contours resulting from multi-frame relaxation are shown in Figure 4c. Note that 3D filtering operations have the effect of blurring the features into the sixth frame. For comparison, the result obtained with our earlier proximal contour tracking method [14] is shown in Figure 4d. The statistical surface tracking method again performs well. It is interesting to note that it still tracks effectively over the most inclined portion of the sphere close to the pole. Here the radial velocity of the contours is 4 pixels per frame. By contrast, the proximal method loses contact at a much lower latitude when the radial contour velocity exceeds 2 pixels per frame. Moreover, the proximal method has a tendency to pursue features into the void beyond the physical extent of the sphere. The statistical surface tracking method cleanly disengages once the minimum depth on the sphere is reached.

In order to explore the capacity of our method to accommodate changes in contour topology, we have generated a synthetic volumetric image in which two cylinders merge to form a single cylinder. Viewed at the slice level the sequence represents one of two topological events depending on the tracking direction. In one direction two closed contour merge together to form a single closed contour. In the opposite direction, a single closed contour breaks into two fragments. Figure 5a shows the fragmentation event, with the segmented result below.

5 Conclusions

We have described a novel approach to surface tracking in volumetric image stacks. The method uses a statistical model of the uncertainties inherent in the

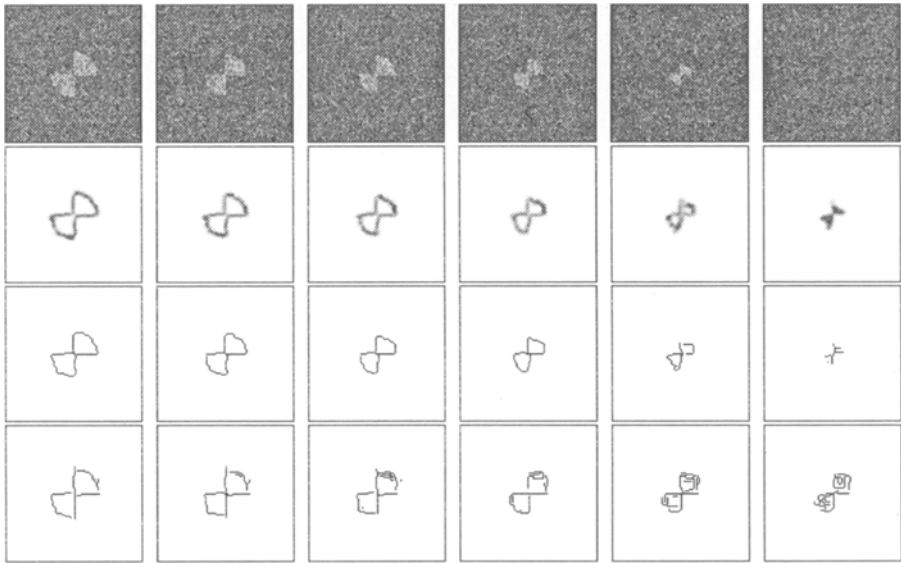


Fig. 4. Top (a) Sequence of synthetic images; Second row (b) Statistical contour field; Third row (c) Labeled contours; Bottom row (d) Labeled contours with proximal contour tracking.




Fig. 5. Top (a) Synthetic sequence; Bottom (b) Labeled contours.

characterisation of intensity surfaces to compute an evidential field for inter-frame contour displacements. This field is obtained using Gaussian density kernels which are parameterised in terms of the variance-covariance matrices for predicted contour displacements. The underlying variance model accommodates both the effects of raw image noise on estimated surface normals. The evidential field effectively couples contour displacements to the intensity features on successive frames through a process of statistical contour tracking. Hard contours are extracted using a dictionary-based relaxation process.

References

1. Blake A., Curwen, R. and Zisserman A., "Affine Invariant Contour Tracking with Automatic control of Spatiotemporal Scale", *Proceedings of the Fourth International Conference on Computer Vision*, pp. 66-75, 1993.

2. Bruckstein A.M. and Shaked D., "Projective Invariant Smoothing and Evolution of Planar Curves", *Aspects of Visual Form Processing*, Edited by C. Arcelli, L. Cordella and G. Sanniti di Baja, pp. 109–119, 1994.
3. Gueziec A., "Large Deformable Splines, Crest Lines and Matching", *Proceedings of the Fourth International Conference on Computer Vision*, pp 650–657, 1993.
4. Hancock E.R. and Kittler J., "Combining Evidence in probabilistic relaxation", *International Journal of Pattern Recognition and Artificial Intelligence*, **3**, pp.29–52, 1989. .
5. Hancock E.R. and Kittler J., "Edge labelling using dictionary based relaxation", *IEEE PAMI*, **PAMI 12**, pp. 165–181, 1990.
6. Hancock E.R. and Kittler J., "Adaptive estimation of hysteresis thresholds", *Proceedings IEEE Computer Vision and Pattern Recognition Conference*, pp. 196–201, 1991.
7. Hancock E.R., "Resolving Edge-Line Ambiguities using Probabilistic Relaxation", *Proceedings IEEE CVPR Conference*, pp. 300–306, 1993.
8. Ivins J. and Porrill J., "Active Region Models for Segmenting Texture and Colour", *Image and Vision Computing*, **13**, pp. 431–438, 1995.
9. Kass M., Witken A. and Terzopoulos D., "Constraints on deformable models: Recovering 3D Shape from non-rigid motion", *Artificial Intelligence*,**36**, pp 91–123, 1988 .
10. Kimia B.B., Tannenbaum A. and Zucker S.W., "On the Shape Triangle", *Aspects of Visual Form Processing*, Edited by C. Arcelli, L. Cordella and G. Sanniti di Baja, pp. 307–323, 1994.
11. Monga O., Deriche R., and Malandain G., "Recursive Filtering and edge closing; two primary tools for 3D edge detection", *Image and Vision Computing*, **9**, pp. , 1991.
12. Osher S.J. and Sethian J.A., "Fronts Propagating with Curvature Dependent Speed: Algorithms Based on Hamilton-Jacobi Formulations", *Journal of Computational Physics*, **79**, pp. 12–49, 1988.
13. Sander P.T. and Zucker S.W., "Inferring surface trace and differential structure from 3D images", *IEEE PAMI*, **PAMI 12**, pp. 833–854, 1990.
14. Sharp N.G. and Hancock E.R., "Feature Tracking by Multi-frame Relaxation", *Image and Vision Computing*, **13**, pp. 637–644, 1995.
15. Zucker S.W, David C., Dobbins A., and Iverson L., "The organisation of curve detection: tangent fields and fine-spline coverings", *Proc. Second Int. Conf. Computer Vision*, pp. 577–586, 1988.