

# Tracking Elongated Structures using Statistical Snakes

Ricardo Toledo, Xavier Orriols, Xavier Binefa, Petia Radeva, Jordi Vitrià and J.J. Villanueva \*

Computer Vision Center and Dpt. d'Informàtica

Universitat Autònoma de Barcelona

Edifici O, 08193, Bellaterra, SPAIN

{ricardo, xevi, xavierb, petia, jordi, juanjo}@cvc.uab.es

## Abstract

*In this paper we introduce a statistic snake that learns and tracks image features by means of statistic learning techniques. Using probabilistic principal component analysis a feature description is obtained from a training set of object profiles. In our approach a sound statistical model is introduced to define a likelihood estimate of the grey-level local image profiles together with their local orientation. This likelihood estimate allows to define a probabilistic potential field of the snake where the elastic curve deforms to maximise the overall probability of detecting learned image features. To improve the convergence of snake deformation, we enhance the likelihood map by a physics-based model simulating a dipole-dipole interaction. A new extended local coherent interaction is introduced defined in terms of extended structure tensor of the image to give priority to parallel coherence vectors.*

**Keywords.** Deformable Models, Snakes, Statistic Learning, PPCA, Medical Imaging.

## 1. Introduction

Given the problem of tracking flexible elongated objects in images, we focus on dynamic contours (snakes) [6, 5, 7, 9, 3]. Snakes have the advantage to allow to incorporate high-level knowledge into the image processing in terms of different constraints such as smoothness, continuity [6], guide by an approximate model [8], etc. These constraints regularize the problem of image feature organization and guide the process of finding a solution as a function of the initial conditions. Snakes [6, 3, 8, 9] are dynamic elastic curves with a physical interpretation. Energy principles are used to deform the elastic curve. Usually, feature map is generated as an edge/crest/valley map [6] and the snake deforms on a potential field constructed as a distance map to the extracted image features. Therefore, the snake is highly affected by the quality of a specific image feature detector. We consider statistic learning of linear

structures to cope with the appearance variance, the search space is reduced by adding knowledge from the application domain. Obtaining better response from feature detectors the global snake performance is increased. In [7, 3] Principal Component Analysis (PCA) is used to model face shape and grey-level images. A limiting disadvantage of PCA is the absence of a probability density model and an associated likelihood measure. The need of a probability density framework is clearly present in problems where *saliency* is formulated in terms of visual similarity. The derivation of PCA from a perspective of density estimation offers the advantage that the probability density function gives a measure of the novelty of a new data point [2].

Given the advantages of Probabilistic PCA (PPCA) as a straightforward technique to construct statistic image feature description [2] and snakes as a global segmentation and tracking technique [6], in this paper we propose a combination of these techniques to track non-rigid elongated objects. Object profiles are learned from a training set and a statistic classifier is constructed as a potential field of the snake. The map is built in two steps: first, a structure tensor is applied to assign a coherence direction to each pixel of the target image. Second, image profiles perpendicular to the coherence orientation are weighed by PPCA defining a likelihood map. Therefore, each point has assigned a probability measure to belong to the learned feature category. Additionally, we refine the likelihood map applying to the coherence direction field an extended local coherent detection between neighbours of the likelihood map as a function of derivatives up to the second order. We show that this detection prioritises regions of pixels with parallel coherent direction improving the localisation of elongated structures. As a result, the refined likelihood map has good responses around the object while small amount of false responses are observed. To avoid false stationary states of the snakes we introduce a hybrid potential map as a combination of refined likelihood map and distance map that assures slow movement far from the object and fast convergence when approaching the object of interest.

\*This work was supported by CICYT and EU grants TAP98-0631, TEL99-1206-C02-02, TIC98-1100 and 2FD97-0220.

The outline of the paper is as follows: section 2 summarizes the snake technique and the PPCA. Section 3 focuses on the statistic framework for feature learning. Section 4 presents the detection process of linear structures. Section 5 and 6 demonstrate the feasibility to improve the statistic map adding more specific domain knowledge. The article finishes with results and conclusions.

## 2. Fundamentals

The problem of linear structure tracking is solved by two main techniques: snakes for a global segmentation approach and PPCA to obtain a robust image feature detection. This section briefly introduces these basic concepts.

### 2.1. Snakes

Snakes represent physics-based models defined from the theory of elasticity and Newton Mechanics. Elastic properties of physical nature as tension and stiffness are attributed to the deformable models to control its deformation. The snake is deforming as a physical body where the deformation is smooth and without discontinuities.

The segmentation by snakes is defined as an energy minimization problem. The snake deforms minimizing its external energy (i.e. approaching as close as possible to the image features of interest), and at the same time minimizes its internal energy keeping its shape as smooth as possible. Representing the position of the snake parametrically as  $u(s) = (x(s), y(s))$ , the energy functional of the snake can be written as:

$$E_{snake} = \int E_{int}(u(s)) + E_{ext}(u(s)) ds$$

where  $u(s)$  is the elastic curve,  $E_{int}(u(s))$  is the internal energy, and  $E_{ext}(u(s))$  the external (image) energy.

External forces push the snake towards the image features of interest. These forces are associated to a potential  $P(x, y)$  which, in general, is defined in terms of gradient module of the image convolved by a Gaussian function [6] or as a distance map of the edge points:

$$E_{ext} = P(x, y) = d(\vec{x}), \quad P(\vec{x}) = -e^{-d(\vec{x})^2} \quad (1)$$

where  $d(\vec{x})$  denotes the distance between pixel  $\vec{x}$  and its closest edge. The snake is moved by potential forces and tries to fall in a valley as if it was under the effect of gravity.

Following the general energy-minimizing scheme, we replace the deterministic potential map of the snake by a new hybrid potential map. To obtain this statistic map the features of interest characterizing the object model are learned and the snake potential is built as a measure of the probabilities of finding the learned feature in every point of

the image under analysis. Deforming the snake maximizes the probability to approach image features that best represent the object corresponding to the learned model.

### 2.2. Probabilistic Principal Component Analysis

In the PPCA framework [2] a small number of causes are considered, that acting in combination generate the complexity of the observed data set. This leads to define a joint distribution  $p(\vec{t}, \vec{x})$  over visible  $\{\vec{t}\}$  and hidden variables  $\{\vec{x}\}$ , the corresponding distribution  $p(\vec{t})$  for the observed data is obtained by marginalization:

$$p(\vec{t}) = \int p(\vec{t} | \vec{x}) p(\vec{x}) d\vec{x}$$

The main goal is to find the parameters that maximize the joint observed data distribution i.e. the best description under a specific generative model.

One of the basic tools is the standard factor analysis [1], which seeks to relate  $d$ -dimensional observed data vectors  $\{\vec{t}_n\}$  corresponding to a set of  $q$ -dimensional latent variables  $\{\vec{x}_n\}$  by a linear mapping:


$$\vec{t}_n = W \vec{x}_n + \vec{\mu} + \vec{\epsilon} \quad (2)$$

where latent variables are distributed into an isotropic Gaussian distribution,  $\sim \mathcal{N}(\vec{0}, I)$ . The noise model  $\vec{\epsilon}$ , or error, is considered also Gaussian such that  $\vec{\epsilon} \sim \mathcal{N}(\vec{0}, \Psi)$ , the  $(d \times q)$  parameter matrix  $W$  contains the factors loading, and  $\mu$  is a constant which, maximized the likelihood, corresponds to the mean  $\vec{\mu}$  of the data. Given this formulation, the model for  $\vec{t}$  is also normal  $\mathcal{N}(\vec{\mu}, \Sigma)$ , with mean  $\vec{\mu}$  and covariance matrix  $\Sigma = WW^T + \Psi$ .

Assuming uniformly distributed noise over the whole image and linearity assumption in (2) lead to the development of a PPCA [2]. In this case  $\Psi$  endows with equal variance the principal axes (i.e.  $\Psi = \sigma^2 I$ ). Hence, PPCA is a permissible technique when illuminant variations problem is not analyzed from variance structure. Considering this key assumption leads to consider the *conditional independence* of observed data. The underlying idea is that the dependencies between data variables  $\vec{t}$  are explained by a small number of latent variables  $\vec{x}$ , while  $\vec{\epsilon}$  represents the unique variance of each observation variable. Instead, conventional PCA treats both variance and covariance identically. The corresponding distribution of observed data ( $D = \{\vec{t}_1, \dots, \vec{t}_N\} / \vec{t}_n \in \mathbb{R}^d$ ) defines the Mahalanobis distance, in terms of log-likelihood:

$$\mathcal{L}(\vec{\mu}, \Sigma) = -\log p(D | \vec{\mu}, \Sigma) \quad (3)$$

At this point, the problem is centered on parameter estimation, which, in practice, will be given by data observations. This leads to consider the problem of *incomplete*



**Figure 1.** Learning of profiles perpendicular to the vessel direction (a) and a profile from the training set (b).

*data.* For this purpose, Dempster et al. in [4] use the EM algorithm, where each observation  $\vec{t}_n$  is associated to an unobserved state  $\vec{x}_n$ , and the main goal is to determine which component generates the observation. In this sense, the unobserved states can be seen as *missing data* and therefore the union of observations  $\vec{t}_n$  and  $\vec{x}_n$  is said to be complete data,  $\vec{y}_n = (\vec{t}_n, \vec{x}_n)$ . In this way the likelihood measure to be maximized is the *Complete-log-Likelihood*, i.e.  $\mathcal{L} = \sum_{n=1}^N \log p(t_n, x_n)$ . Maximum-likelihood formulation for PPCA also allows a closed solution for the mapping matrix  $W$  and the noise variance  $\sigma$  [2]:

$$W = U_q (\Lambda_q - \sigma^2 I)^{\frac{1}{2}} R; \quad \sigma^2 = \frac{1}{d-q} \sum_{j=q+1}^d \lambda_j \quad (4)$$

where  $U_q$  are the first  $q$  eigenvectors of the data set covariance matrix,  $\Lambda_q$  is a diagonal matrix with the corresponding first  $q$  eigenvalues ( $\lambda_i, \forall 1 \leq i \leq N$ ) and  $R$  is an arbitrary rotation matrix.

### 3. A Probabilistic Framework for Structure Learning

To build the statistic map of the snake, a learning process is carried out from a training set of profiles of the elongated object in images. We define a descriptor representation by local grey-level image profiles perpendicular to the object elongation and longer than the maximal expected linear structure width (fig. 1).

To obtain the samples, we use a "classic" snake [6] that deforms on a potential generated by a crease based distance map and corrected by the user. Each profile is defined by an image coordinate point  $(i, j)$  at the center of the structure and a direction  $\vec{v}$  that indicates its orientation. The learning process is done by choosing the central points  $(i, j)$  and orientation of profiles  $\vec{v}$  from the snake position.

Let  $U, V$  be vectorial spaces referred to the image coordinates  $(i, j)$  and to the orientations  $\vec{v}$  respectively, where  $(i, j) \in U$  and  $\vec{v} \in V$ . Then, we define a mapping

$\mathcal{T} : U \times V \rightarrow \mathbb{R}^d$ , relating each extended coordinate pair and orientation  $(i, j; \vec{v}) \in U \times V$ , to a grey-level profile  $\vec{t}_n \in \mathbb{R}^d$ :  $\vec{t}_n = \mathcal{T}(i, j; \vec{v})$

By sampling the training set of images with a profile length  $d$  enough to cover the widest structure, we obtain a learning data set  $D = \{\vec{t}_1, \dots, \vec{t}_N\}$ . All profiles are aligned in order to obtain as far as possible axial symmetry with respect to the middle structure point. An interesting consequence of such an alignment is that the maximal probability shall be assigned to the real crease of the elongated object.

Actually, linear structure profiles are characterized by a high degree of statistic regularity due to their morphological consistency. Ought to the high number of pixels (samples) in any image, a dimensional space reduction by means of PPCA is used to statistically describe the feature. Each sample is a grey-level profile and the covariance matrix  $S$  of the observed data is constructed from the learning data set:  $S = \frac{1}{N} \sum_{n=1}^N (\vec{t}_n - \vec{\mu})(\vec{t}_n - \vec{\mu})^T$ , where  $N$  corresponds to the profile population,  $\mu$  is the sample mean profile and  $\vec{t}_n$  is the  $n$ -th grey-level profile. The diagonalization of the covariance matrix  $S$  allows to reach the closed solution for the maximum likelihood estimation in (4). Hence, it provides the transformation between latent variables and observed data as a linear mapping  $\mathbb{R}^q \rightarrow \mathbb{R}^d$ , being  $q$  the latent space dimension and  $d$  the profile length, defined by the projection matrix  $W$  and the sample mean  $\vec{\mu}$ .

### 4. Feature detection


At this stage, the main goal is to build a local likelihood map for the snake given the image neighbourhoods projected onto the optimized space  $\mathbb{R}^d$ . The mapping is a probability measure of belonging to the learned structure category  $\Omega = (\Sigma, \vec{\mu})$  for each coordinate pixel  $(i, j)$  and any direction  $\vec{v}$ . Since the learned model is done through the gray-level profiles, we define the following functional composition to build the likelihood map  $\mathcal{P}(\mathcal{T}(i, j; \vec{v}) | \Omega)$ :

$$\begin{aligned} U \times V &\rightarrow \mathbb{R}^d \rightarrow \mathbb{R} \\ (i, j; \vec{v}) &\rightarrow \vec{t} \rightarrow \mathcal{P}(\mathcal{T}(i, j; \vec{v}) | \Omega) \end{aligned}$$

Computing this probability requires a factorization of the model covariance matrix:  $\Sigma^{-1} = (WW^T + \sigma^2 I)^{-1}$ . This can be done with low computational cost ( $\mathcal{O}(q^3)$ ) instead of ( $\mathcal{O}(d^3)$ ) using the Woodbury's identity:

$$(WW^T + \sigma^2 I)^{-1} = \{I - W(W^T W + \sigma^2 I)^{-1} W^T\} / \sigma^2$$

In order to detect the learned feature, the image under analysis ought to be scanned searching profiles like the learned ones. Hence, the main problem arises when it comes to assign for each pixel  $(i, j)$  the profile orientation  $\vec{v}$ . Instead, we orient the profiles in the direction of maximal



**Figure 2.** Eigenvectors associated with the lowest (a) and largest (b) eigenvalue.

grey-level variance of image neighbourhood by a structure tensor [10]. The structure tensor field, applied to an integration region  $\rho$ , of the regularized gradient image ( $\nabla I_\sigma$ ), under a suitable scale  $\sigma$ , measures the coherence between the regions and the searched structure [10]:

$$J_\rho(i, j) = K_\rho * (\nabla I_\sigma \nabla I_\sigma^T)(i, j) \quad (\rho \geq \sigma \geq 0) \quad (5)$$

where  $(i, j)$  are the image coordinates and  $K_\rho$  is a gaussian convolution kernel (zero mean, variance  $\rho$ ). The eigenvalues  $\mu_{1,2}$  of the tensor (5) ( $\mu_1 \geq \mu_2$ ) describe the average contrast variation in the eigendirections  $\vec{w}_{1,2}$  ( $\vec{w}_1 \perp \vec{w}_2$ ). The eigenvector associated to the lower eigenvalue  $\vec{w}_2$  is the orientation of lowest fluctuation, detecting the elongated flow, fig. 2(a). The first eigenvector describes the directions of maximal grey-level variance, e.g. the direction coincident with the one used to learn the profiles, fig. 2(b). This fact encourages to constrain the degrees of freedom of orientation profile vectors  $\vec{v}$ , since they can be searched from the vector field  $\vec{w}_1$  with the biggest structural tensor eigenvalue (fig.2 (b)).

A normalized coherence measure of local image structure is obtained as follows [10]:  $\kappa = \frac{(\mu_1 - \mu_2)^2}{(\mu_1 + \mu_2)^2}$ . The discontinuity in case of  $\mu_1 = \mu_2 = 0$  is not a problem in our case because the equality is true only in areas where the image is flat. The normalized coherence measure allows to focus on regions of interest with significant value of the coherence measure  $\kappa$  reducing, in this way, the computational cost of generation of the likelihood map. Since profile orientation  $\vec{v}$  is fixed for each pixel  $(i, j)$ , the profile extraction is done through the mapping:  $(i, j; \vec{v}) \rightarrow \vec{t} = \mathcal{T}(i, j; \vec{v})$ . Then the distance map from every profile to the learned model is obtained as a result of assigning its corresponding negative log-likelihood defined in (3) to each pixel (fig. 4). The minimal probability computed in the regions of interest is assigned to all areas discarded by the threshold on the coherence measure  $\kappa$ . In the fig. 3 bright regions correspond to the most probable crease points.

## 5. Extended Local Coherent Interactions

The use of a likelihood-based detection method allows for additional restrictions adapted to specific problems. To provide a robust probabilistic potential to the snake segmentation, the likelihood map is improved taking advantage of local coherence vector field. Notice that coherence directions are parallel in the regions of linear structures (fig. 2). Analogously to the ferromagnetism in solid-state physics, such property induces to consider the vector field as a lattice with interacting dipoles. Each dipole is associated with an orientation vector  $\vec{v}$  provided by the principal eigenvector of the structure tensor in the image pixel  $(i, j)$ . The energy distribution has to stand out two factors: local parallelism and the likelihood distribution. To reward parallelism we take the interaction between vectors as a functional of inner product between neighbour vectors. The weight provided by the likelihood can be seen as the charge associated to each dipole.

A simple energy function  $\mathcal{H}$  of interaction assigned between two points  $\vec{x} = (i, j) \in U$  and  $\vec{y} = (i', j') \in U$  is proportional to the inner product between their associated directions:  $\vec{v}(\vec{x}) \in V$  placed in  $\vec{x}$  and  $\vec{v}(\vec{y}) \in V$  at  $\vec{y}$ :

$$\mathcal{H}(\vec{x}, \vec{y} | \Omega) = \left\{ \mathcal{P}[\vec{x}; \vec{v}(\vec{x}) | \Omega] \mathcal{P}[\vec{y}; \vec{v}(\vec{y}) | \Omega] \right\} \vec{v}(\vec{x})^T \vec{v}(\vec{y}) \quad (6)$$

where  $\mathcal{P}[\vec{x}; \vec{v}(\vec{x}) | \Omega]$  is the probability assigned to the profile centered at  $\vec{x}$  with orientation  $\vec{v}(\vec{x})$  consistent with the model  $\Omega = (\Sigma, \vec{\mu})$ . The interaction energy integrated on a neighbourhood domain  $R$  for each point gives the energy distribution map:


$$H(\vec{x} | \Omega) = \int_R | \mathcal{H}(\vec{x}, \vec{y} | \Omega) | f(|\vec{x} - \vec{y}|) d\vec{y} \quad (7)$$

where  $f(|\vec{x} - \vec{y}|)$  is a decreasing function of the distance between neighbours:

$$f(|\vec{x} - \vec{y}|) = \frac{1}{2\pi\sigma^2} \exp \left\{ -\frac{|\vec{x} - \vec{y}|^2}{2\sigma^2} \right\}$$

and  $\sigma$  is a scale parameter which determines the interaction range. In fact,  $H(\vec{x} | \Omega)$  stands out the relations of coherence between coherent vectors like a *meta-coherence* turning regions of non parallel vectors into homogeneous areas (fig. 3). As a result of applying the local coherence interacting model, the refined likelihood map has local maxima at object centrelines (fig. 3).

Furthermore, we extend the structure tensor to refine the detection of elongated structures expanding (5) by second image derivatives. Therefore, the interacting energy in (6) is automatically extended to four dimensions. The obtained filter behaves more selective in distinguishing image features from other false coherence structures. Fig. 3 shows the




**Figure 3.** Comparison of local coherence interaction (a) with its topographic map (b) and the extended coherence interaction (c) with the corresponding topographic map (d).

effect of including second derivatives into the ferromagnetic interaction dipole-dipole model. On the right-low corner one can see the selective filtering of false positive responses of the approach. Note that a combination of first and second derivatives into the interaction energy differentiates to greater extent elongated "valley-like" structures from other objects in the image (e.g. the vessel and the diaphragm that can be observed in the original image from fig. 4(a)). In this example, diaphragm part gives lower response by the extended coherence approach due to the use of the first and second derivatives in the extended structure tensor.

## 6. Hybrid Snake Potential

The obtained likelihood map offers high responses at centrelines of elongated objects and nearby form strong slope towards the centreline (fig. 3). As a result, good convergence behaviour is observed once the snake falls in a neighbourhood of the elongated structure. Another advantage of the refined map is the small amount of false responses. Therefore, far from the objects of interest the potential map consists of large plateaus with constant 0 likelihood. In these regions the snake initialised does not suffer any external forces to deform it. To cope with it, we enhance the negative likelihood map  $\hat{p}(\vec{x}|\Omega) = -p(\vec{x}|\Omega)$  obtaining a hybrid potential field as follows:

$$P_{hybrid}(x|\Omega) = \begin{cases} \hat{p}(\vec{x}), & \hat{p}(\vec{x}) \in [-1, 0) \\ d(\vec{x}), & \text{otherwise.} \end{cases}$$




**Figure 4.** Seeds for initial snake (a), likelihood map (b), hybrid probabilistic map (c) and topographic map (d) of a part of the hybrid potential shown by the black square in (c)

where  $d(\vec{x})$  is the distance to the closest potential point with non-zero probability. As a result, the snake moves with a constant small step towards the potential zones and converges only to the zones corresponding to learned image features. Fig. 4 illustrates an original image, the likelihood map according to (4), the hybrid map and a topographic map corresponding to the black square in fig. 4(c).

## 7. Experimental Results

We build and test the hybrid potential maps tracking a coronary tree vessel in a sequence of angiographies. The vessels are dynamic elastic elongated objects in images with very low contrast and high signal-to-noise rate. A more sophisticated approach to segment and track objects is obligatory in order to achieve good results. For each image frame  $I$  in the sequence, a likelihood map is built and the minimized snake in frame  $I-1$  is used as initialization. The first snake initialization is obtained using a path search through the vessel under analysis. The user provides the starting point of the path and the path search is carried out as follows: from the starting point the maximum likelihood vector is chosen and a set of coherence vectors is searched in the neighborhood for maximum scalar product. The process is repeated up to the end of the path (fig. 4(a)).

To build the maps a previous learning step is made with



**Figure 5.** Tracking by statistic snakes

a set of twenty image frames. We take two hundred samples from each image and each sample is of forty pixels width. The sampled profiles are obtained perpendicular to the centerline of the vessels. To obtain the centerline, we use a conventional snake and the perpendicular profiles are defined regarding the snake. For each learned vessel we obtain a matrix where the rows are the profiles along a vessel (fig. 1). We make the dimensional reduction through PPCA using the first five principal axes, which explain up to 97% of the observed data variance in feature space.

The built map comprises the extraction of coherent directions through the second moment matrix (fig. 2) discarding regions with low module of the coherent vectors to speed up the process of map generation. Afterwards, the profiles over the extracted directions are obtained using as the middle point of the profile the origin of each direction vector. The profiles length, as in the learning step, is forty pixels. The profiles are then compared against the learned ones and as a result a measure of the probability of being a true vessel profile is obtained. Using more specific knowledge about the domain, we recover the coherence directions and after weighting them with the probabilities we test for parallelisms. Non parallel vectors mean high probability of being a false profile. As shown in (7), there is a compromise between the region of integration  $R$  and the scale  $\sigma$  related to homogeneity and vessel thinness precision. We used as integration region  $R$  a square of  $25 \times 25$  pixels and a scale  $\sigma = 1.5$ . Fig. 5 illustrates the snake tracking for a set of consecutive frames.

## 8. Summary and Conclusions

The value of the statistic basis for linear structure detection and tracking has been established by demonstrating the mechanism of the PPCA embedded into the snake framework. In order to manage complex objects and the variability of appearance of image structures, our technique is supported by a learning approach to extract and detect only these "crease-like" features determined by the training set. Learned models are used in a probabilistic framework in order to build significant energy potential, resulting in less false responses of the image feature detector and more robust snake-based object tracking.

A new approach to potential computation using a likelihood map is formulated and applied to the tracking of specific structure on angiographies: coronary vessels as a difficult example of automatic analysis. The initialization process demands only one point to the user achieved by exploiting the coherence vector field and the likelihood function. Finally, the likelihood map is modified to avoid the uncertainty regions far from the detected feature using a distance function to the high likelihood regions. As a result, the snake is less dependent on its initialisation and once placed on the hybrid potential map it converges to image features with high probability to represent learned object profiles. The obtained results and the self-training capability of the snake encourage utilizing it in different applications.

## References

- [1] D. Bartholomew. *Latent Variable Models and Factor Analysis*. Charles Griffin and Co. Ltd., 1987.
- [2] C. M. Bishop and M. E. Tipping. Mixtures of probabilistic principal component analysers. *Tech. Report NCGR*, 1998.
- [3] A. Blake and M. Isard. *Active Contours*. Spr. Verlag, 1998.
- [4] A. Dempster, N. Lair, and D. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society Series B*, 39:1–38, 1977.
- [5] G. J. Edwards, C. J. Taylor, and T. F. Cootes. Learning to identify and track faces in image sequences. In *British Machine Vision Conference*, volume 8, pages 130–139, 1997.
- [6] M. Kass, A. Witkin, and D. Terzopoulos. Snakes : Active contour models. In *ICCV*, pages 259–268, 1987.
- [7] A. Lanitis, C. J. Taylor, and T. F. Cootes. A unified approach to coding and interpreting face images. In *ICCV*, volume 5, pages 368–373, 1995.
- [8] P. Radeva, J. Serrat, and E. Martí. A snake model-based segmentation. In *ICCV*, pages 816–821, 1995.
- [9] R. Toledo, P. Radeva, C. von Land, and J. Villanueva. 3D dynamic model of the coronary tree. In *Computers in Cardiology*, pages 777–780, 1998.
- [10] J. Weickert. Coherence-enhancing diffusion of colour images. In *Image and Vision Computing*, volume 17, 1999.