Modeling Bayesian Estimation for Deformable Contours

In Proceedings of 7th IEEE International Conference on Computer Vision. Kerkyra, Greece. September 20-27, 1999.

Stan. Z. Li Juwei Lu
School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore
szli@szli.eee.ntu.edu.sg http://markov.eee.ntu.edu.sg:8000/˜szli/

Abstract

A novel trainable snake model, called EigenSnake, is presented in the Bayesian framework. In the EigenSnake, prior knowledge of a specific object shape, such as that of face outlines and facial features, is derived from a training set of the shape and incorporated into a Bayesian snake model in the form of the prior distribution. Further, a "shape space", which is constructed on the basis of a set of eigenvectors obtained from principle component analysis, is used to restrict and stabilize the search for the optimal solution. The effectiveness is demonstrated by experiments, which shows that the EigenSnake produces more reliable and accurate results than existing models.

1 Introduction

Correct location and extraction of objects of interest is an important step in many applications such as industrial inspection, medical image analysis, and object recognition. Deformable models, such as snake [7], G-Snake [8], Active Shape Model (ASM) [2] and deformable templates [6], are able to fit objects more closely and have proven to be more effective in detecting shapes subject to local deformation. Such models can be classified into two classes [6, 5]: (i) free-form, and (ii) parametric. In the free-form class, such as the snakes [13, 14], there is no global structure of the template; the template is constrained by only local continuity and smoothness constraints. Such models can be deformed to any arbitrary shapes deviating greatly from the shape of interest.

In the parametric class, a global shape model is designed and encoded by a small number of parameters [16, 4, 12, 8, 2, 11, 6, 1]. *Prior* knowledge about the structural properties is incorporated into the shape models, assuming that prior information of the geometrical shape or a set of training samples is available. Compared with the free-form class, the parametric class is more robust against to irrelevant structures, occlusion and bad initialization. A quite successful and versatile scheme in this class is that using statistical shape models in the Bayesian framework [3, 12, 11, 8, 6, 1]. In these models, both the prior knowledge and observation statistics are used to define the optimal, Bayesian estimate.

In this paper, a novel snake model, called EigenSnake, is presented in the Bayesian framework (Section 2). Three strategies are combined into the EigenSnake to make the object seeking more robust and accurate: (i) The prior knowledge of a specific object shape is derived from a training set of shape samples, and expressed as a prior distribution for defining a Bayesian estimate as the optimal solution for the object contour. The prior distribution can explain and represent the global and local variation in the training set. (ii) Constraints on object shapes are imposed in a "soft" manner, following the derived prior distribution. The shape encoding can adjust itself dynamically based on the up-to-date knowledge learned during the solution finding process. This adjustment allows the solution to fit more closely to the object shape. (iii) A shape space is constructed based on a set of eigenvectors derived by principle component analysis (PCA) performed on the training set. It is used to restrict and stabilize the search for the Bayesian optimum.

The affine invariant *shape matrix* method in the G-Snake [8] is used for encoding the global shape information. However, we noticed a drawback of the original method: It encodes the shape information in a "hard" manner in that the shape matrix is fixed. This actually assumes that the considered object shape undergoes only *rigid* (such as affine) transformation. So, only a small amount of local deformation is tolerated, and an incorrect solution results when the shape deformation is large, which is often a problem in many applications such as face detection. Because of this, a better solution, which is a contour fitting more closely to the object in the image, can have a higher energy value than an inferior solution. In contrast, in the EigenSnake, the shape matrix is dynamically modifiable based on the up-to-date information gathered from the solution finding process. This helps produce a more accurate result.

The shape space constraint used in the ASM [2, 9] is applied onto the solution finding process. The shape space, which is a subspace of the whole arbitrary contour space, explains and represents most of the non-rigid variance in the training set, based on the eigenvectors derived through PCA performed on the training contours. Every intermediate solutions (energy minima) is projected into the subspace and thus the solution is always restricted to be within the subspace. This proves to be effective in stabilizing the solution. However, in the ASM, no priori distribution accounting for global rigid motion is incorporated; therefore, the original ASM is unable to handle large scale and orientational shape-changes. This is amended in the EigenSnake.

Experiments are performed in the application domain of face extraction (Section 3). Four methods based on the G-Snake and the EigenSnake are compared to evaluate how various constraints and different prior knowledges have affected the results. Results demonstrate that the EigenSnake, as compared to the other three methods, is more stable to changes in scale and rotation and more robust to spurious structures and occlusions. It is also shown that the EigenSnake effectively overcomes problems due to initialization, occlusion and noise and produces much more accurate results.

2 Formulation of EigenSnake

The EigenSnake incorporates the prior knowledge about the distribution of training contours. The global and local variations in the training set are well represented by the resulting prior distribution. Moreover, a shape space is constructed to further constrain the solution.

2.1 Joint Prior Distribution

Let $\mathbb{E} = \{(x, y) \in \mathbb{R}^2\}$ be the image plane. An arbitrary contour can be represented by a vector of m points in $\mathbb{E}$, $f = [f_1^T, f_2^T, \cdots, f_m^T]$ (a $2 \times m$ matrix) where $f_i \in \mathbb{E}$ is a point. Given a prototype contour $\overline{f}$, a shape matrix $A = A(\overline{f})$ can be determined by $A(\overline{f})\overline{f}^T = \mathbf{0}$. The shape matrix is under affine transformation: If f^{aff} is an affine-transformed version of $\overline{f}$, then $A(\overline{f}){f^{aff}}^T = \mathbf{0}$ [8].

For the convenience of the following description, we express any contour as a $2m \times 1$ vector $f = [f_1, f_2, \cdots, f_m]^T \in \mathbb{R}^{2m}$. Given a prototype contour $\overline{f}$ and the corresponding shape matrix $A(\overline{f})$, the prior distribution of f is given as

$$p(f \mid \overline{f}) \propto \exp\left\{-E(f \mid \overline{f})\right\} \quad (1)$$

where $E(f \mid \overline{f}) = \sum_{i=1}^{m} \frac{E(f_i \mid \overline{f})}{\sigma_i^2}$ is the internal energy [8].

In the G-Snake, $\overline{f}$ is chosen to be the mean of the training contours and is fixed. This actually assumes that the considered object shape (that of $\overline{f}$) is subject to *rigid* (such as affine) transformation only. However, it ignores prior knowledge about the fact that local variations exist among the training contours. This is equivalent to the assumption that f is a distribution with zero variance, which is obviously invalid with the training set. Since $\overline{f}$, and hence $A(\overline{f})$, is fixed, it is difficult for the solution f to converge closely to the object contour in the image when the local shape deformation is large (this is often the case for many applications such as face detection). Because of this, a better solution, which is a contour fitting more closely to the object in the image, can have a higher energy value than an inferior solution, as will be illustrated by experimental results. This is referred to as the incorrectness of an energy model.

The EigenSnake amends this problem by allowing $\overline{f}$ to deform according to a prior distribution derived from the training contour. In this case, the following joint prior is considered

$$P(f, \overline{f}) = p(f \mid \overline{f}) \cdot p(\overline{f}) \propto \exp\{-E(f, \overline{f})\} \quad (2)$$

Hence, the shape information is encoded in a "soft" manner, and $\overline{f}$ and hence $A(\overline{f})$ can be modified dynamically during the computation. The prior $p(\overline{f})$ is derived from a set of training contours in the following.

2.2 Learning Prior from Training Set

Let a set of L face contours $\{f^1, f^2, \ldots, f^L\}$ be given as the training set where each contour f^i is represented by a sequence of m evenly placed (done manually) points in $\mathbb{E}$ (*i.e.* $f_k^i \in \mathbb{E}$ and $f^i \in \mathbb{R}^{2m}$). The training contours are aligned at the centers of their gravities, and then scaled and rotated by using an iterative alignment algorithm so as to remove the rigid variation in the training samples as much as possible. Thus pre-processed contours constitute a random distribution. The mean contour $\overline{f}_{mean}$ is computed by averaging the corresponding contour points. See Fig.1.

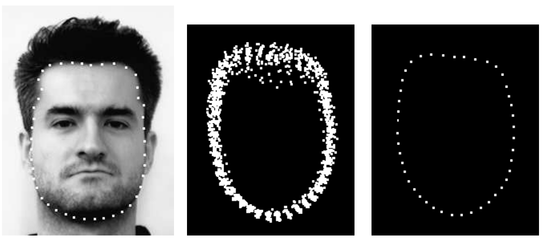

Figure 1. Left: A sample of face contour with $m = 44$ points. Middle: Scatter of points from normalized and aligned set of face boundary shapes. Right: The mean (average) of the training face contours.

Assume that the training samples obey a Gaussian distribution with the mean vector $\overline{f}_{mean}$ and the covariance matrix C:

$$p(f) \propto \exp\left\{-\frac{1}{2}\left(f-\overline{f}_{mean}\right)^T C^{-1}\left(f-\overline{f}_{mean}\right)\right\} \tag{3}$$

The distribution represents the prior knowledge about local variation derived from the training contours. However, the covariance matrix C is often singular, thus C^{-1} non-existing, because the dimensionality of the samples is high and the number of available training contours is often small. The distribution has to be approximated by some means.

2.3 Approximation of Prior Using PCA

We adopt an optimal method proposed by Moghaddam and Pentland [10] to approximate the Gaussian density $p(f)$ by applying a PCA to the training samples to capture the main modes of the variation in the directionally elongated (*i.e.* elliptical) Gaussian distribution.

Let Ψ be the eigenvector matrix of C, and Λ be the corresponding diagonal matrix of eigenvalues. We have $\Lambda = \Psi^T C \Psi$. The first $M (\ll 2m)$ eigenvectors are selected: $\Phi = [\phi_1, \phi_2, \cdots, \phi_M]$. They are used as the basis vectors to constitute a specific *shape space* (*i.e.* the M-dimensional space spanned by the M orthonormal vectors). The shape space efficiently explains most of the local variation within the training set. Any contour $f \in \mathbb{R}^{2m}$ can be approximated by a linear combination of the M eigenvectors [15]

$$\hat{f} = \overline{f}_{mean} + \Phi\omega \tag{4}$$

where $\omega = \Phi^T(f - \overline{f}_{mean})$ is the vector of weights. $\hat{f}$ is the projection of f into the shape space, and ω is the coordinate of the projection point in the shape space. It provides an optimum approximation for f in the sense of least squares error. It discards the variation of f which is inconsistent with those reflected by training contours, while retaining most of the variation that the shape space can explain. The projection $\hat{f}$ has the following distribution [10]

$$p(\hat{f}) = p(\omega) = \frac{\exp\left\{-\frac{1}{2}\sum_{k=1}^{M}\frac{\omega_k^2}{e_k}\right\}}{(2\pi)^{M/2}\prod_{k=1}^{M} e_k^{1/2}} \tag{5}$$

where e_k is the k-th largest eigenvalue of C, $\omega_k = \omega_k(f) = \phi_k^T\left(f - \overline{f}_{mean}\right)$, $\varepsilon^2(f) = \left\|f - \overline{f}_{mean}\right\|^2 - \sum_{k=1}^{M}\omega_k^2$ is the residual error.

2.4 Prior Distribution of Prototype Contour

During solution finding process, the prototype contour $\overline{f}$ is modified dynamically according to the up to date MAP solution, and is confined in the shape space. The method is described below.

Denote the MAP solution currently found at iteration t by $f' = \arg\min_{f\in\mathcal{N}(f^{(t-1)})} E(f, \overline{f} \mid d)$ where $E(f, \overline{f} \mid d)$ is the posterior energy (to be defined later) and $\mathcal{N}(f^{(t-1)})$ is the neighborhood of previously found $f^{(t-1)}$. The MAP estimate f' is aligned with $\overline{f}_{mean}$ by minimizing $\left\|u + s \cdot R_\theta(f') - \overline{f}_{mean}\right\|$ where u (translation), s (scaling) and θ (rotation) are the rigid transformation parameters and $R_\theta(f)$ is a function that rotates the contour f by θ. The translation vector can be found as $u = (x_u, y_u) = \frac{1}{m}\sum_{i=1}^{m}[\overline{f}_{mean,i} - f'_i]$. The other two parameters can be calculated using a simple alignment algorithm. The aligned contour is given as $f'' = u + s \cdot R_\theta(f')$.

The aligned contour f'' is then projected into the shape space using Equ.(4), resulting in the projected version $\hat{f}''$. The prototype contour $\overline{f}$ is set to

$$\overline{f} = \hat{f}'' = \overline{f}_{mean} + \Phi\omega \tag{6}$$

where $\omega = \Phi^T(f'' - \overline{f}_{mean})$ is the coordinates of f'' in the shape space.

Therefore, $\overline{f}$ obtained in such a way follows the same distribution as that of $\hat{f}''$

$$p(\overline{f}) = p(\hat{f}'') = p(\omega) \tag{7}$$

as in Equ.(5). The corresponding prototype contour energy is

$$E(\overline{f}) = \frac{1}{2}\sum_{i=1}^{M}\frac{\omega_i^2}{e_i} \tag{8}$$

Note that $p(f \mid \overline{f})$ is scale and rotation invariant, and $p(\overline{f})$ depends only on the local deformation parameter vector ω and is independent of the global deformation parameter such as u, s, θ. As a consequence, $p(f, \overline{f})$ is also scale and rotation invariant, retaining the advantage of the prior distribution of the G-Snake.

To take the advantage of the shape space constraint, the solution found in iteration t is correspondingly revised into

$$f^{(t)} = -u + \frac{1}{s} \cdot R_{-\theta}(\hat{f}'') \tag{9}$$

The revised solution $f^{(t)}$ is in the shape space. It retains the rigid variation of f' (w.r.t. $\overline{f}_{mean}$), and has the local deformation consistent with the samples in the training set.

Two remarks about the imposition of the shape space constraints follow: (i) The imposition of the shape space constraint makes the EigenSnake more robust against spurious structures and occlusion. The effect will be demonstrated later in the experiments section. (ii) Because $\omega = \mathbf{0}$ when $f'' = \overline{f}_{mean}$, ω can be considered as the local deformation parameters of f'' or f'. Hence, the rigid transformation parameters (u, s and θ) and the non-rigid deformation parameters (ω) are separated. Therefore, the prior distribution $p(f, \overline{f})$ explains not only the global variation but also the local variation, derived from a training set.

2.5 Likelihood Distribution

Edges are detected from the image $d = \{d(x, y)\}$. The normalized edge direction at each pixel location (x, y), denoted $d^e(x, y) = (d_x^e(x, y), d_y^e(x, y))$, is used to provide constraints on the data, with $\|d^e(x, y)\| \in (0, 1]$ if an edge is detected at (x, y) or a zero vector $d^e(x, y) = \mathbf{0}$ if an edge is not detected there.

Denote the unit tangent directions of f by $f^e = [f_1^e, \ldots, f_m^e]$. The direction at a point $f_i = (x_i, y_i)$ can be estimated by using a finite difference formula: $\frac{f_{i+1}-f_i}{|f_{i+1}-f_i|} + \frac{f_i - f_{i-1}}{|f_i - f_{i-1}|}$. It is normalized into a unit-length vector: $f_i^e = (f_{x\,i}^e, f_{y\,i}^e)$, with $\|f_i^e(x, y)\| = 1$.

Constraints on the image d are imposed through the differences $\|f_i^e - d_i^e\|$. Assuming the differences have the same variance σ_e^2 for all i, we define the likelihood energy function as

$$E(d \mid f) = \frac{1}{\sigma_e^2} \sum_{i=1}^{m} \|f_i^e - d_i^e\|^2 \tag{10}$$

where $d_i^e = d^e(x_i + x_u, y_i + y_u)$. The likelihood distribution can be defined as

$$p(d \mid f) = \frac{1}{Z'} \exp\{-E(d \mid f)\} \tag{11}$$

where Z' is the normalizing constant

2.6 Bayesian Estimates

In the Bayesian framework, the problem of extracting a contour with unknown deformation from a given image can be posed as a problem of maximum a posterior (MAP) estimation. The posterior distribution of the EigenSnake is then defined as

$$p(f, \overline{f} \mid d) = \frac{p(d \mid f) \cdot p(f, \overline{f})}{p(d)} \tag{12}$$

where $p(d \mid f) = p(d \mid f, \overline{f})$. The corresponding energy is

$$E(f, \overline{f} \mid d) = E(d \mid f) + E(f \mid \overline{f}) + E(\overline{f}) \tag{13}$$

The MAP estimates f_{MAP} and $\overline{f}_{MAP}$ are then defined as

$$\{f_{MAP}, \overline{f}_{MAP}\} = \arg\min_{f, \overline{f}} \{E(f, \overline{f} \mid d)\} \tag{14}$$

2.7 MultiEigenSnake

The EigenSnake is expanded into the MultiEigenSnake for detecting a compound object with multiple separate sub-objects. Suppose a compound object consists of L simple sub-objects. The shape of each sub-object is represented by a simple contour, and their combination is represented by a compound contour. Therefore, the MultiEigenSnake consists of $L + 1$ contours. For example, a

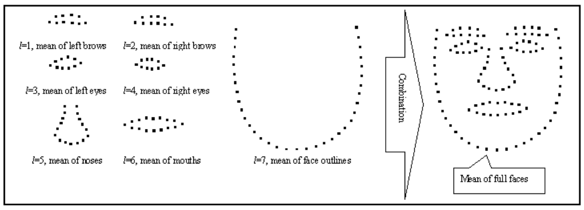

Figure 2. Means of training contours for MultiEigenSnake.

face, is considered as composed of $L = 7$ sub-objects: a face outline, a pair of brows, a pair of eyes, a nose and a mouth. Fig.2 shows the means of training contours and their combination for the representation of a full face.

Let $\overline{f^{ME}}$ denotes the prototype contours of the compound object, and $\overline{f}_\ell$ denotes the prototype contour of the ℓ-th sub-object. Corresponding to Equ.(13), we define the following posterior energy for the MultiEigenSnake

$$E(\overline{f^{ME}}, f^{ME}, f^1, \cdots, f^L, \overline{f^1}, \cdots, \overline{f^L} \mid d) = E(\overline{f^{ME}}) + \sum_{\ell=1}^{L} [E(d \mid f^\ell) + E(f^\ell \mid \overline{f^\ell})] \tag{15}$$

The prior energy $E(\overline{f^{ME}})$ accounts for prior knowledge about local deformation of the compound object and the spatial configuration of its L sub-objects.

3 Experiments

Two sets of experiments are presented. The first examines the effects of using the specific prior knowledge and the shape space constraint by comparing the EigenSnake and other methods in extracting face outlines. The second set demonstrates the efficacy of the MultiEigenSnake in the extraction of the full facial features.

3.1 Comparison EigenSnake with Other Methods

Here, four methods are compared: the original G-Snake (GS), enhanced G-Snake with the shape space constraint imposed (GS+), EigenSnake without the imposition of the shape space (ES-), and the full EigenSnake (ES). Their respective performances will tell us how the incorporation of the specific prior knowledge (by comparing ES- vs GS, and ES vs GS+) and the imposition of the shape space constraint help improve the extraction of face outline (ES vs ES-, and GS+ vs GS).

The experiment data set contains 60 frontal face images of size 512×342 from the face database of Bern university. $m = 44$ landmark points are manually and evenly placed on the face outline of each training image as accurately as possible. A subset of 40 images are used as the *training set* and all the 60 images are used as the *testing set*.

The mean contours, which was shown in Fig.1:Right, is calculated from the training set. 8 eigenvectors are used to form the shape space. $M = 35$ eigenvectors are used to estimate the prior energy $E(\overline{f})$.

Contour initialization is done manually, as a common practice for snakes. It has the same shape as the mean contour but it is re-scaled randomly to between 70% and 120% of the true mean scale, shifted away randomly from the true position by between ± 20 pixels, and rotated randomly from the true orientation in the range of $\pm 15°$. All the methods have the same initial contour for each test image.

Table 1. Accuracy of the four methods.

Method	Init	GS	GS+	ES-	ES
$D(\text{method})$	23.8	17.1	12.7	7.7	4.8

Table 1 compares the accuracy of face outlines detected by the four methods (plus that of the initial outline), where the accuracy is measured in terms of the average Euclidean distance (in pixels), denoted by $D(\text{method})$, between corresponding points in the extracted outline and the manual marked outline. The compared results illustrate that both the prior energy $E(\overline{f})$ and the shape space constraint can improve the accuracy. The improvement in accuracy made by the use of the shape space constraint may be calculated as $[(D(\text{GS}) - D(\text{GS+})) + (D(\text{ES}-) - D(\text{ES}))]/2 = 3.65$ (the average of the two improvements, one from GS to GS+, the other from ES- to ES). The improvement in accuracy made by the use of the prior energy $E(\overline{f})$ may be calculated as $[(D(\text{GS}) - D(\text{ES}-)) + (D(\text{GS+}) - D(\text{ES}))]/2 = 8.65$ (the average of the two improvements, one from GS to ES-, the other from GS+ to ES). The two figures suggest that the contribution of the prior energy $E(\overline{f})$ is greater than that of the shape space constraint; in other words, more job was done by the imposition of the specific prior knowledge than by the use of the shape space constraint.

Fig.3 shows an example of face outline extraction. It can been seen that the GS often fails to localize objects. Although GS+ can modify the results of GS so that it looks more like "face outline", the improvement is less than that produced by ES- and ES. The two ES methods, which use the new energy functions with abundant prior knowledges incorporated, can effectively resist the disturbance from spurious structures, noisy edge map and improper initial scales.

The final energy values tell us the "correctness" of the four energy-minimization based models (for a correct energy model, a better solution should have a lower energy). The GS+ produced better results than the GS; however, the final GS+ energies are higher than corresponding GS energies. This means that the GS energy does not really define the optimal solution in its minimum and is therefore "incorrect". This is due to lack of constraint on the optimal solution. In contrast, the ES model has a proper energy function because a better solution is always associated with a lower energy.

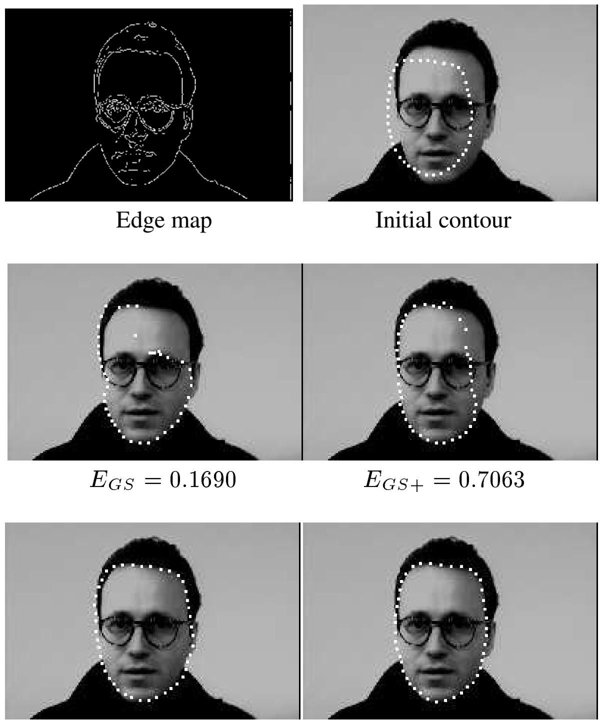

Edge map | Initial contour

$E_{GS} = 0.1690$ | $E_{GS+} = 0.7063$

$E_{ES-} = 0.8258$ | $E_{ES} = 0.7216$

Figure 3. Comparison of the four methods.

3.2 Extracting Full Face Using MultiEigenSnake

This experiment demonstrates the MultiEigenSnake applied to extract a full face with $L = 7$ facial features. The training set consists of 120 frontal face images from the face database of Bern university. The testing set is chosen from the face database of Carnegie Mellon University.

A total of 108 landmark points is used to represent a full face, among which 10 for left brow, 10 for right brow, 10 for left eye, 10 for right eye, 17 for nose, 16 for mouth and 35 for the face boundary. The means of the training contours (shown as Fig.2) are used as initial prototype contours of L sub-models and the compound model respectively.

Some experimental results (initial (left) and final contours)are shown in Fig.4. Basically, the MultiEigenSnake can successfully extract the full face contours if the initial distance between the initial prototype contour and the target object is less than one fourth of the initial prototype

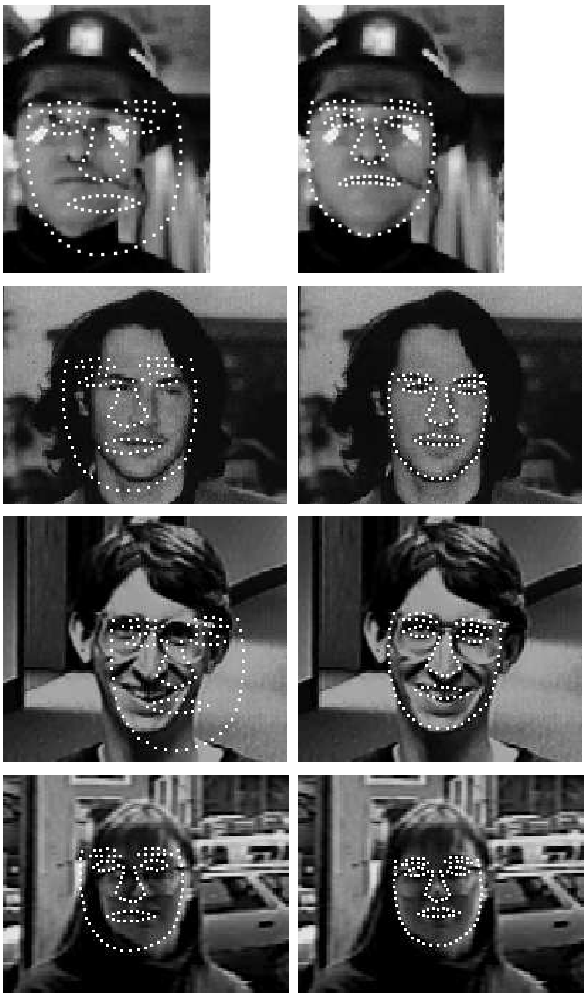

Figure 4. MultiEigenSnake results.

contour size. The MultiEigenSnake has certain scale insensitive and rotation insensitive; for example, it allows up to $\pm 20°$ rotation of initial prototype contour from orientation of the true object, as shown in Fig.4.

4 Conclusion

In this paper, we have presented a novel Bayesian model, EigenSnake, for extraction of deformable contours. The EigenSnake is trainable, general and capable of representing an arbitrary shape. It incorporates prior knowledge about both rigid and non-rigid deformations of expected shapes, and deforms to fit the contour in a way which reflects variations in the training set. The EigenSnake has the ability not only to account for global changes due to rigid motions, but also to learn local deformation of contours from examples. Experimental results show that the EigenSnake is more robust to initialization, spurious features, and occlusion, and can produce more accurate contours than existing methods. The EigenSnake may be applied to other applications such as motion tracking. A more comprehensive report of this work can be found at http://markov.eee.ntu.edu.sg:8000/~szli/papers/iccv99.ext.ps.

Acknowledgement This work was supported by NTU-AcRF RG 43/95 and RG 51/97.

References

[1] A. Blake, B. Bascle, M. Isard, and J. MacCormick. “Statistical models of visual shape and motion”. *Phil. Trans. R. Soc. Lond. A*, 356:1283–1302, 1998.

[2] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. “Active shape models: Their training and application”. *CVGIP: Image Understanding*, 61:38–59, 1995.

[3] U. Grenander. *Pattern synthesis: Lectures in pattern theory*. Springer-Verlag, 1976.

[4] U. Grenander and D. M. K. Y. Chow. *Hands : a pattern theoretic study of biological shapes*. Springer-Verlag, New York, 1991.

[5] A. K. Jain, Y. Zhong, and M.-P. Dubuisson-Jolly. “Deformable template models: A review”. *Signal Processing*, 71:109–129, 1998.

[6] A. K. Jain, Y. Zhong, and S. Lakshmanan. “Object matching using deformable templates”. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(3):267–278, March 1996.

[7] M. Kass, A. Witkin, and D. Terzopoulos. “Snakes: Active contour models”. *International Journal of Computer Vision*, 1:321–331, 1988.

[8] K. F. Lai and R. T. Chin. “Deformable contour: modeling and extraction”. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17:1084–1090, 1995.

[9] A. Lanitis, C. J. Taylor, and T. F. Cootes. “Automatic interpretation and coding of face images using flexible models”. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19:743–756, 1997.

[10] B. Moghaddam and A. P. Pentland. “Probabilistic visual learning for object representation”. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19:696–710, 1997.

[11] D. Mumford. “Pattern theory: a unified perspective”. In D. Knill and W. Richard, editors, *Perception as Bayesian Inference*, pages 25–62. Cambridge University Press, 1996.

[12] L. Staib and J. Duncan. “Boundary finding with parametrically deformable medels”. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14:1061–1075, 1992.

[13] D. Terzopolous, J. Platt, A. Barr, and K. Fleischer. “Elastically deformable models”. *Comput. Graphic*, 21(4):205–214, 1987.

[14] D. Terzopolous, A. Witkin, and M. Kass. “Constraints on deformable models: Recovering 3d shape and nonrigid motion”. *AI*, 36:91–123, 1988.

[15] M. A. Turk and A. P. Pentland. “Eigenfaces for recognition”. *Journal of Cognitive Neuroscience*, 3(1):71–86, March 1991.

[16] A. L. Yuille, D. Cohen, and P. W. Hallinan. “Feature extraction from faces using deformable templates”. In *Proc. CVPR*, pages 104–109, 1989.