Modelling Human Shape with Articulated Shape Mixtures

Abdullah A. Al-Shaher and Edwin R. Hancock

University of York, York YO1 5DD, UK

Abstract. This paper describes a statistical framework for recognising 2D shapes with articulated components. The shapes are represented using both geometrical and a symbolic primitives, that are encapsulated in a two layer hierarchical architecture. Each primitive is modelled so as to allow a degree of articulated freedom using a polar point distribution model that captures how the primitive movement varies over a training set. Each segment is assigned a symbolic label to distinguish its identity, and the overall shape is represented by a configuration of labels. We demonstrate how both the point-distribution model and the symbolic labels can be combined to perform recognition using a probabilistic hierarchical algorithm. This involves recovering the parameters of the point distribution model that minimise an alignment error, and recovering symbol configurations that minimise a structural error. We apply the recognition method to human moving skeleton.

1 Introduction

The task of recognising articulated shapes has attracted considerable interest in computer vision. The main problem is how to robustly recover correspondence when the object being tracked undergoes deformations and the detected feature points defining the object are subject to noise. One of the most effective ways of developing matching techniques is to draw on probabilistic and statistical methods. This approach has lead to the development of point distribution models [1], deformable templates [2] and condensation [3].

There are a number of ways in which object articulation can be modelled. Perhaps the simplest of these is to decompose the shape into a skeletal form, consisting of limbs or branches, and to model the articulation of the branches. The mechanical movement of the resulting shape can be captured by the rotation of the components. However, in order to constrain the change in shape to be physically realistic bounds, or distributions, must be imposed on the rotation angles [4, 5]. Hence, the mechanical constraints on articulation must be combined with a statistical model of limb movement. In addition to movement of the limbs, the articulated shape also has a structural composition, since the limbs can be assigned labels to distinguish them, and the arrangement of the labels can be used to provide further constraints for shape-recognition.

The aim in this paper is to develop a statistical framework that can be used to recognise articulated shapes using information concerning limb movement and

symbolic constraints concerning the overall shape structure. To do this, we develop a hierarchical algorithm. Each shape is represented as an arrangement of articulated limbs. The movement of the limbs is represented by a polar point distribution model. The structural component of the model is represented by a configuration of limb-labels. The recognition architecture has two intercommunicating layers. The first of these is concerned with limb alignment, and this aims to recover the lengths and polar angles of the limbs. The second aims to assign limb-labels so that the overall structure is recovered consistently.

2 Point Distribution Models

The point distribution model of Cootes and Taylor commences from a set training patterns. Each training pattern is a configuration of labelled point co-ordinates or landmarks. The landmark patterns are collected as the the object in question undergoes representative changes in shape. To be more formal, each landmark pattern consists of L labelled points whose co-ordinates are represented by the set of position co-ordinates $\{X_1, X_2, \dots, X_L\} = \{(x_1, y_1), \dots, (x_L, y_L)\}$. Suppose that there are T landmark patterns. The t^{th} training pattern is represented using the long-vector of landmark co-ordinates $X_t = (x_1, y_1, x_2, y_2, \dots, x_L, y_L)^T$, where the subscripts of the co-ordinates are the landmark labels. For each training pattern the labelled landmarks are identically ordered. The mean landmark pattern is represented by the average long-vector of co-ordinates $Y = \frac{1}{T} \sum_{t=1}^T X_t$. The covariance matrix for the landmark positions is

$$\Sigma = \frac{1}{T} \sum_{t=1}^T (X_t - Y)(X_t - Y)^T \quad (1)$$

The eigenmodes of the landmark covariance matrix are used to construct the point-distribution model. First, the unit eigenvalues E of the landmark covariance matrix are found by solving the eigenvalue equation $|\Sigma - EI| = 0$ where I is the $2L \times 2L$ identity matrix. The eigen-vector ϕ_i corresponding to the eigenvalue E_i is found by solving the eigenvector equation $\Sigma \phi_i = E_i \phi_i$. According to Cootes and Taylor [1], the landmark points are allowed to undergo displacements relative to the mean-shape in directions defined by the eigenvectors of the covariance matrix Σ . To compute the set of possible displacement directions, the M most significant eigenvectors are ordered according to the magnitudes of their corresponding eigenvalues to form the matrix of column-vectors $\Phi = (\phi_1 | \phi_2 | \dots | \phi_M)$, where $E_1, E_2, \dots, E_M$ is the order of the magnitudes of the eigenvectors. The landmark points are allowed to move in a direction which is a linear combination of the eigenvectors. The updated landmark positions are given by $\hat{X} = Y + \Phi \gamma$, where γ is a vector of modal co-efficients. This vector represents the free-parameters of the global shape-model. When fitted to an observed set of landmark measurements X_o , the least-squares estimate of the parameter vector is

$$\gamma = \frac{1}{2} (\Phi + \Phi^T)(X_o - Y)$$

3 Shape Representation

Our aim is to use point distribution models to account for shape deformations due to limb articulation. The model is a two component one. First, we have a limb-model. This accounts for the variations in shape of each of the individual limbs using a point distribution model to describe the modes of variation of the landmark points about a mean shape. Second, we have a limb arrangement model. This is an augmented point distribution model that describes the arrangement of the centre points of the limbs, and their polar angles.

We are concerned with recognising 2D shapes by modelling segment movement around the centre of the shape. The shape under study is assumed to be segmented into a set of K jointed and non-overlapping limbs. The k^{th} limb is represented by a long-vector of consecutive landmark points

$$X_k = (x_1^k, y_1^k, x_2^k, y_2^k, \dots, x_{n_k}^k, y_{n_k}^k)^T$$

The centre-of-gravity of the limb indexed k is

$$c_k = \frac{1}{|n_k|} \sum_{i=1}^{n_k} (x_i^k, y_i^k)^T$$

The overall shape is represented by a long-vector of consecutive limb centres $C = (c_1^T, c_2^T, \dots, c_K^T)^T$. The centre of articulated shape is computed by averaging the centre of the limbs

$$U = \frac{1}{|K|} \sum_{k=1}^K c_k$$

To model the articulated shape, we use a polar variant of the standard point distribution model [6]. This model allows the primitives to move about the centre of articulation. According to this model the shape is viewed as an arrangement of non-overlapping primitives. Each primitive is represented by mean point c_k . Limb articulation is represented by a set of limb-angles. For the k^{th} limb the angle is defined to be

$$\theta_k = \tan^{-1} \frac{U(y) - c_k(y)}{U(x) - c_k(x)}$$

and the angular arrangement of the limbs is represented by the vector $\Theta = (\theta_1, \theta_2, \dots, \theta_K)^T$. The global movement of the limbs within a shape is specified by the concatenated long-vector of angles and the centres-of-articulation, i.e. by the vector $S = (\Theta^T, C^T)^T$.

To augment the geometric information, we assign symbols to the articulated components. Each training pattern is assigned to a shape class and each component primitive is assigned to a primitive class. The set of shape-labels is Ω_c and the set of articulated component or limb labels is Ω_s . The symbolic structure of each shape is represented a permissible arrangement of limb-labels. For shapes of class $\omega \in \Omega_c$ the permissible arrangement of limbs is denoted by $\Lambda_\omega = \langle \lambda_1^\omega, \lambda_2^\omega, \dots \rangle$.

4 Learning Mixtures of PDM's

In Cootes and Taylor's method [7], learning involves extracting a single covariance matrix from the sets of landmark points. Hence, the method can only reproduce variations in shape which can be represented as linear deformations of the point positions. To reproduce more complex variations in shape either a non-linear deformation or a series of local piecewise linear deformations must be employed.

In this paper we adopt an approach based on mixtures of point-distributions. Our reasons for adopting this approach are twofold. First, we would like to be able to model more complex deformations by using multiple modes of shape deformation. The need to do this may arise in a number of situations. The first of these is when the set of training patterns contains examples from different classes of shape. In other words, we are confronted with an unsupervised learning problem and need to estimate both the mean shape and the modes of variation for each class of object. The second situation is where the shape variations in the training data can not be captured by a single covariance matrix, and a mixture is required.

Our approach is based on fitting a Gaussian model to the set of training examples. We commence by assuming that the individual examples in the training set are conditionally independent of one-another. We further assume that the training data can be represented by a set of shape-classes Ω . Each shape-class $\omega \in \Omega_s$ has its own mean point-pattern Y_ω and covariance matrix Σ_ω . With these ingredients, the likelihood function for the set of training patterns is

$$p(X_t, t = 1, \dots, T) = \prod_{t=1}^T \sum_{\omega \in \Omega_s} p(X_t | Y_\omega, \Sigma_\omega) \quad (2)$$

where $p(X_t | Y_\omega, \Sigma_\omega)$ is the probability distribution for drawing the training pattern X_t from the shape-class ω . According to the EM algorithm, we can maximise the likelihood function above, by adopting a two-step iterative process. The process revolves around the expected log-likelihood function

$$Q_{n+1} = \sum_{t=1}^T \sum_{\omega \in \Omega_s} P(t \in \omega | X_t, Y_\omega^{(n)}, \Sigma_\omega^{(n)}) \ln p(X_t | Y_\omega^{(n+1)}, \Sigma_\omega^{(n+1)}) \quad (3)$$

where $Y_\omega^{(n)}$ and $\Sigma_\omega^{(n)}$ are the estimates of the mean pattern-vector and the covariance matrix for class ω at iteration n of the algorithm. The quantity $P(t \in \omega | X_t, Y_\omega^{(n)}, \Sigma_\omega^{(n)})$ is the *a posteriori* probability that the training pattern X_t belongs to the class ω at iteration n of the algorithm. The probability density for the pattern-vectors associated with the shape-class ω , specified by the estimates of the mean and covariance at iteration $n+1$ is $p(X_t | Y_\omega^{(n+1)}, \Sigma_\omega^{(n+1)})$. In the M, or maximisation, step of the algorithm the aim is to find revised estimates of the mean pattern-vector and covariance matrix which maximise the expected log-likelihood function. The update equations depend on the adopted model for the class-conditional probability distributions for the pattern-vectors.

In the E, or expectation, step the *a posteriori* class membership probabilities are updated. This is done by applying the Bayes formula to the class-conditional density. At iteration $n + 1$, the revised estimate is

$$P(t \in \omega | X_t, Y_\omega^{(n)}, \Sigma_\omega^{(n)}) = \frac{p(X_t | Y_\omega^{(n)}, \Sigma_\omega^{(n)}) \pi_{t,\omega}^{(n)}}{\sum_{\omega \in \Omega} p(X_t | Y_\omega^{(n)}, \Sigma_\omega^{(n)}) \pi_{t,\omega}^{(n)}} \quad (4)$$

where

$$\pi_{t,\omega}^{(n+1)} = \frac{1}{T} \sum_{t=1}^T P(t \in \omega | X_t, Y_\omega^{(n)}, \Sigma_\omega^{(n)}) \quad (5)$$

4.1 Mixtures of Gaussians

We now consider the case when the class conditional density for the training patterns is Gaussian. Here we assume that the pattern vectors are distributed according to the distribution

$$p(X_t | Y_\omega^{(n)}, \Sigma_\omega^{(n)}) = \frac{1}{(2\pi)^L \sqrt{|\Sigma_\omega^{(n)}|}} \exp \left[-\frac{1}{2} (X_t - Y_\omega^{(n)})^T (\Sigma_\omega^{(n)})^{-1} (X_t - Y_\omega^{(n)}) \right] \quad (6)$$

At iteration $n + 1$ of the EM algorithm the revised estimate of the mean pattern vector for class ω is

$$Y_\omega^{(n+1)} = \sum_{t=1}^T P(t \in \omega | X_t, Y_\omega^{(n)}, \Sigma_\omega^{(n)}) X_t \quad (7)$$

while the revised estimate of the covariance matrix is

$$\Sigma_\omega^{(n+1)} = \sum_{t=1}^T P(t \in \omega | X_t, Y_\omega^{(n)}, \Sigma_\omega^{(n)}) (X_t - Y_\omega^{(n)}) (X_t - Y_\omega^{(n)})^T \quad (8)$$

When the algorithm has converged, then the point-distribution models for the different classes may be constructed off-line using the procedure outlined in Section 2.

We apply this learning procedure separately to the landmark data for the individual limbs, and to the combined limb angle and limb-centre data. For the limb with label λ , the estimated modal matrix is Φ_λ and the estimated parameter vector is γ_λ . For the shape-class with label ω , on the other hand, the combined modal matrix for the articulation angles and limb-centres is $\tilde{\Phi}_\omega$ and the result of fitting to data is a parameter vector $\tilde{\Gamma}_\omega$. The first K rows of $\tilde{\Gamma}_\omega$ correspond to the limb angles, and the remaining $2K$ to the long-vectors of limbs centres. However, we need to constrain the parameters corresponding to the limb angles. Suppose that the mean-vector for the limb-angles is $\tilde{\Theta}_w$ and the corresponding covariance matrix is Σ_w . The angular deformations are constrained to avoid flipping by limiting the deformation vector. We use the variance associated with the eigenmodes to constrain the deformation. The k^{th} component of the parameter vector

is constrained to fall in the interval $-3\sqrt{E_k} \leq \Gamma_{(k)} \leq 3\sqrt{E_k}$. The articulation angles lie in the range -180° to 180° to avoid discontinuities associated with the flip from 0° to 360° . A similar procedure for learning is followed to learn the variation in the polar representation of the limb and limb classes.

5 Hierarchical Architecture

With the limb-articulation and limb-centre point distribution models to hand, our recognition method proceeds in a hierarchical manner. Our aim is to classify the set of limb landmark long-vectors $X = \{z_1, \dots, z_k, \dots, z_K\}$ representing a test-shape. To commence, we make maximum likelihood estimates of the best-fit parameters of each limb-model to each set of limb-points. The best-fit parameters γ_λ^k of the limb-model with class-label λ to the set of points constituting the limb indexed k is

$$\gamma_\lambda^k = \arg \max_{\gamma} p(z_k | \Phi_\lambda, \gamma) \quad (9)$$

We use the best-fit parameters to assign a label to each limb. The label is that which has maximum a posteriori probability given the limb parameters. The label assigned to the limb indexed k is

$$l_k = \arg \max_{l \in \Omega_s} P(l | z_k, \gamma_\lambda, \Phi_\lambda) \quad (10)$$

In practice, we assume that the fit error residuals follow a Gaussian distribution. As a result, the class label is that associated with the minimum squared error. This process is repeated for each limb in turn. The class identity of the set of limbs is summarised by the string of assigned limb-labels $L = \langle l_1, l_2, \dots \rangle$. Hence, the input layer is initialised using maximum likelihood limb parameters and maximum a posteriori probability limb labels. The shape-layer takes this information as input. The goal of computation in this second layer is to refine the configuration of limb labels using global constraints on the arrangement of limbs to form consistent shapes. The constraints come from both geometric and symbolic sources. The geometric constraints are provided by the fit of a polar limbs point distribution model. The symbolic constraints are provided by a dictionary of permissible limb-label strings for different shapes.

The parameters of the limb-centre point distribution model are found using the EM algorithm [8]. Here we borrow ideas from the hierarchical mixture of experts algorithm [9], and pose the recovery of parameters as that of maximising a gated expected log-likelihood function for the distribution of limb-centre alignment errors $p(X | \Phi_\omega, \Gamma_\omega)$. The likelihood function is gated by two sets of probabilities. The first of these are the a posteriori probabilities $P(\lambda_k^\omega | z_k, \gamma_{\lambda_k^\omega}, \Phi_{\lambda_k^\omega})$ of the individual limbs. The second are the conditional probabilities $P(L | A_\omega)$ of the assigned limb-label string given the dictionary of permissible configurations for shapes of class ω . The expected log-likelihood function is given by

$$\mathcal{L} = \sum_{\omega \in \Omega_c} P(L | A_\omega) \left\{ \prod_k P(\lambda_k^\omega | z_k, \gamma_{\lambda_k^\omega}, \Phi_{\lambda_k^\omega}) \right\} \ln p(X | \tilde{\Phi}_\omega, \Gamma_\omega) \quad (11)$$

The optimal set of polar limb arrangement parameters satisfies the condition

$$\Gamma_{\omega}^* = \arg \max_{\Gamma} P(L|\Lambda_{\omega}) \left\{ \prod_k P(\lambda_k^{\omega} | \mathbf{z}_k, \gamma_{\lambda_k^{\omega}}, \tilde{\Phi}_{\lambda_k^{\omega}}) \right\} \ln p(X|\tilde{\Phi}_{\omega}, \Gamma_{\omega}) \quad (12)$$

From the maximum likelihood alignment parameters we identify the shape-class of maximum *a posteriori* probability. The class is the one for which

$$\omega^* = \arg \max_{\omega \in \Omega_c} P(\omega|X, \tilde{\Phi}_{\omega}, \Gamma_{\omega}^*) \quad (13)$$

The class identity of the maximum *a posteriori* probability shape is passed back to the limb-layer of the architecture. The limb labels can then be refined in the light of the consistent assignments for the limb-label configuration associated with the shape-class ω

$$l_k = \arg \max_{\lambda \in \Omega_s} P(\lambda | \mathbf{z}_k, \gamma_l^k, \tilde{\Phi}_{\lambda}) P(L(\lambda, k) | \Lambda_{\omega}) \quad (14)$$

Finally, the maximum likelihood parameters for the limbs are refined

$$\gamma_k = \arg \max_{\gamma} p(\mathbf{z}_k | \tilde{\Phi}_{l_k}, \gamma, \Gamma_{\omega}^*) \quad (15)$$

These labels are passed to the shape-layer and the process is iterated to convergence.

6 Models

To apply the model to shape-recognition, we require models of the alignment error process and the label error process.

6.1 Alignment Errors

To develop a useful alignment algorithm we require a model for the measurement process. Here we assume that the observed position vectors, i.e. $\mathbf{z}_k$ are derived from the model points through a Gaussian error process. According to our Gaussian model of the alignment errors,

$$p(\mathbf{z}_k | \tilde{\Phi}_{\lambda}, \gamma_{\lambda}) = \frac{1}{2\pi\sigma} \exp \left[-\frac{1}{2\sigma^2} (\mathbf{z}_k - \hat{X}_{\lambda} - \tilde{\Phi}_{\lambda} \gamma_{\lambda})^T (\mathbf{z}_k - \hat{X}_{\lambda} - \tilde{\Phi}_{\lambda} \gamma_{\lambda}) \right] \quad (16)$$

where σ^2 is the variance of the point-position errors which for simplicity are assumed to be isotropic. A similar procedure may be applied to estimate the parameters of the polar limb-angle distribution model.

6.2 Label Assignment

The distribution of label errors is modelled using the method developed by Hancock and Kittler [10]. To measure the degree of error we measure the Hamming distance between the assigned string of labels L and the dictionary item Λ . The Hamming distance is given by

$$H(L, \Lambda_\omega) = \sum_{i=1}^K \delta_{l_i, \lambda_i^\omega} \quad (17)$$

where δ is the Dirac delta function. With the Hamming distance to hand, the probability of the assigned string of labels L given the dictionary item Λ is

$$P(L|\Lambda_\omega) = K_p \exp[-k_p H(L, \Lambda_\omega)] \quad (18)$$

where $K_p = (1-p)^K$ and $k_p = \ln \frac{1-p}{p}$ are constants determined by the label-error probability p .

Fig. 1. Human Training Sets

7 Experiment

We have evaluated our approach on human motion sequences. Figure 1 shows some of the data used for the purpose of learning. In total, we used 14 distinct classes of human motion for learning purposes. We made use of 18 frames per class where each frame is segmented into 13 feature points representing the centre of a limb. In figure 2, we show four example shapes recovered as the output of our learning stage.

To evaluate recognition performance, we have used 1200 frames corresponding to different motion classes for testing. Table 1 shows the label number of

Fig. 2. Learnt Shapes:(a) Boxing, (b) Kicking, (c) Relaxing, (d) Stretching

Table 1. Label replacement for shape classes

Iter. No.	Boxing		Kicking		Relaxing		Picking		Stretching		Shooting	
	Cor	Wro	Cor	Wro	Cor	Wro	Cor	Wro	Cor	Wro	Cor	Wro
1	4	9	2	11	3	10	5	8	4	9	8	5
2	2	11	1	12	2	11	3	10	3	10	10	3
3	1	12	0	13	0	13	2	11	2	12	12	1
4	0	13	0	13	0	13	0	13	0	13	13	0
5	0	13	0	13	0	13	0	13	0	13	13	0

Table 2. Recognition rate for shape classes

Shape	Samples	Correct	Wrong
Boxing	200	198	2
Kicking	200	163	37
Relaxing	200	197	3
Picking	200	181	19
Stretching	200	200	0
Shooting	200	179	21
Recognition Rate		93.16%	6.83%

correct and incorrect limb label assignment as a function of iteration number. It is apparent that the error rate decreases for classes that are irrelevant to the shape under recognition, while it increases for the correct class. Table 2 shows the recognition rate for frames of the six classes of motion. The recognition rate is of 93.16%. The poorest recognition occurs for the kicking, the picking and the shooting classes. Since these classes share similar limb movement, we can conclude that recognition is reasonably high. In Figure 3, we show the alignment process as a function of iteration number. The different curves are for different motion classes. It is clear from the graph that the *a posteriori* probability converges on a clear ambiguous assignment from a state in which there is a considerable ambiguity. Figure 4 shows the recognition rate for different classes when random jitter is present. The recognition rate is reasonably high even under noise.

Fig. 3. Alignment as a function of iteration number

Fig. 4. Recognition Rates with respect to random point movement

8 Conclusion

In this paper, we have described a method for fitting articulated shape-models to landmark point data. The shape deformation process adopted is based on point distribution models. The model representation is a hierarchical one. There is a Cartesian deformation model for the limbs and the limb-centres, together with a polar model which represents limb articulation. We develop a probabilistic framework for fitting a mixture of articulated models to data. The method delivers good results of human shape modelling.

References

1. Cootes T.; Taylor C. Combining point distribution models with shape models based on finite element analysis. *IVC*, 13(5):403–409, 1995.
2. Duta N.; Jain A.; Dubuisson P. Learning 2d shape models. *International Conference on Computer Vision and pattern Recognition*, 2:8–14, 1999.
3. Michael Isard; Andrew Blake. Contour tracking by stochastic propagation of conditional density. In *Proc. ECCV*, pages 343–356, 1996.

4. J Gonzales; J Varona; F Roca; and J Villanueva. aspace: Action space for recognition and synthesis of human actions. *2nd IWAMDO, Spain*, pages 189–200, 2002.
5. James Rehg ; Takeo Kanade. Visual tracking of high dof articulated structures: an application to human hand tracking. *3rd ECCV, Sweden*, pages 35–46, 1994.
6. Heap T.; Hogg D. Extending the point distribution model using polar coordinates. *Image and Vision Computing*, 14:589–599, 1996.
7. Cootes T.; Taylor C. A mixture models for representing shape variation. *Image and Vision Computing*, 17:403–409, 1999.
8. Dempster A.; Laird N.; Rubin D. Maximum likelihood from incomplete data via the em algorithm. *Journal of Royal Statistical Soc. Ser.*, 39:1–38, 1977.
9. Jordan M.; Jacobs R. Hierarchical mixtures of experts and the em algorithm. *Neural Computation*, 6:181–214, 1994.
10. Edwin R. Hancock; Josef Kittler. Edge-labelling using dictionary-based relaxation. *IEEE Transaction on PAMI*, 12(2):165–181, 1990.