Pose Estimation and Tracking Using Multivariate Regression

Arasanathan Thayananthan*   Ramanan Navaratnam*
Björn Stenger◇   Philip H. S. Torr‡   Roberto Cipolla*

*University of Cambridge, Department of Engineering, Cambridge CB2 1PZ, UK
{at315|rn246|cipolla}@eng.cam.ac.uk

◇Toshiba Cambridge Research Laboratory, Cambridge, CB4 0GZ, UK
bjorn@cantab.net

‡Oxford Brookes University, Department of Computing, Wheatley, Oxford
OX33 1HX, UK, philiptorr@brookes.ac.uk

Abstract

This paper presents an extension of the relevance vector machine (RVM) algorithm to multivariate regression. This allows the application to the task of estimating the pose of an articulated object from a single camera. RVMs are used to learn a one-to-many mapping from image features to state space, thereby being able to handle pose ambiguity.

Key words: Regression, Relevance Vector Machines, Tracking, Articulated Motion

This paper considers the problem of estimating the 3D pose of an articulated object such as the human body from a single view. This problem is difficult due to the large number of degrees of freedom and the inherent ambiguities that arise when projecting a 3D structure into the 2D image [5,10]. Once the pose estimation task is solved, temporal information can be used to smooth motion and resolve potential pose ambiguities. This divides continuous pose estimation into two distinct tasks: (1) estimate a distribution of possible configurations from a single frame, (2) combine frame-by-frame estimates to obtain smooth trajectories.

Generative methods estimate the pose by projecting a geometric model into the scene and evaluating a likelihood function that measures agreement with the image. Single frame pose estimation then becomes a complex optimization problem that can be approached with methods such as dynamic programming [8], MCMC [12] or hierarchical search [19].

(a)   (b)

Fig. 1. **(a) Multiple mapping functions.** *Given a single view, the mapping from image features to pose is inherently one-to-many. Mutually exclusive regions in state space can correspond to overlapping regions in feature space. This ambiguity can be resolved by learning several mapping functions from the feature space to different regions of the state space.* **(b) Feature extraction.** *The features are obtained from matching costs (Hausdorff fractions) of shape templates to the edge map. These costs are used for creating the basis function vector* ϕ^{HD}.

Discriminative methods follow a different approach by trying to learn a mapping from image features directly to the 3D pose. A straightforward way to do this is to generate a database of images from a 3D model and use efficient search to find the best match [16,19]. The number of templates required to represent the pose space depends on the range of possible motion and required accuracy, and can be in the order of hundreds of thousands of templates. Only a fraction of the them is searched for each query image, however all templates need to be stored.

The method for hand pose estimation from a single image by Rosales *et al.* addressed some of these issues [15]. Image features were directly mapped to likely hand poses using a set of *specialized mappings*. A 3D model was projected into the image in these hypothesized poses and evaluated using an image based cost function. The features used were low-dimensional vectors of silhouette shape moments, which are often not discriminative enough for precise pose estimation.

Agarwal and Triggs proposed a method for selecting relevant features using RVM regression [1,4]. The image features were shape-contexts descriptors of silhouette points and pose estimation was formulated as a one-to-one mapping from the feature space to pose space. This mapping required about 10% of the training examples. The method was further extended to include dynamic information by joint regression with respect to two variables, the feature vector and a predicted state obtained with a dynamic model [2].

The mapping from silhouette features to state space is inherently one-to-many,

1. Initialize
Partition the training set $\mathcal{V}$ into K subsets by applying the K-means algorithm on the state variable $\mathbf{x}_n$ of each data point v_n. Initialize probability matrix $\mathbf{C}$.

2. Iterate

(i) Estimate regression parameters
Given the matrix $\mathbf{C} \in \mathbb{R}^{N \times K}$, where element $c_{nk} = c_k^{(n)}$ is the probability that sample point n belongs to mapping function k, learn the parameters $\left\{ \mathbf{W}^k, \mathbf{S}^k \right\}$ of each mapping function, by multivariate RVM regression minimizing the following cost function

$$L^k = \sum_{n=1}^{N} c_k^{(n)} \left(\mathbf{y}_k^{(n)} \right)^T \mathbf{S}^k \left(\mathbf{y}_k^{(n)} \right), \text{ where } \mathbf{y}_k^{(n)} = x^{(n)} - \mathbf{W}^k \boldsymbol{\phi}(\mathbf{z}^{(n)}). \quad (1)$$

Note: for speed up, samples with low probabilities may be ignored.

(ii) Estimate probability matrix C
Estimate the probability of each example belonging to each of the mapping function:

$$p(\mathbf{x}^{(n)}|\mathbf{z}^{(n)}, \mathbf{W}^k, \mathbf{S}^k) = \frac{1}{2\pi |\mathbf{S}|^{1/2}} \exp \left\{ -0.5 \left(\mathbf{y}_k^{(n)} \right)^T \mathbf{S}^k \left(\mathbf{y}_k^{(n)} \right) \right\}, \quad (2)$$

$$c_k^{(n)} = \frac{p(\mathbf{x}^{(n)}|\mathbf{z}^{(n)}, \mathbf{W}^k, \mathbf{S}^k)}{\sum_{j=1}^{K} p(\mathbf{x}^{(n)}|\mathbf{z}^{(n)}, \mathbf{W}^j, \mathbf{S}^j)}. \quad (3)$$

Fig. 2. **EM for learning multiple mapping functions $\mathbf{W}_k$**

as similar features can be generated by regions in the parameter space that are far apart, see Figure 1(a). Hence it is important to maintain multiple hypotheses over time. In this paper the pose estimation problem from template matching is formulated as learning one-to-many mapping functions that map from the feature space to the state space. The features are Hausdorff matching scores, which are obtained by matching a set of shape templates to the edge map of the input image, see Figure 1(b). A set of RVM mapping functions is then learned to map these scores to different state-space regions to handle pose ambiguity, see Figure 1(a). Each mapping function achieves sparsity by selecting a small fraction of the total number of templates. However, each RVM function will select a different set of templates. This work is closely related to the work of Sminchisescu et al. [18] and Agarwal et al. [3,4]. Both follow a mixture of experts [13] approach to learn a number of mapping functions (or experts). A gating function is learned for each mapping function during training, and these gating functions are then used to assign the input to one or many mapping functions during the inference stage. In contrast, we use likelihood estimation from projecting the 3D-model to verify the output of each mapping function.

Fig. 3. **RVM regression on a toy dataset.** *The data set consists of 200 samples from three polynomial functions with added Gaussian noise. (a) Initial clustering using K-means. (b), (c),(d) Learned RVM regressors after the 1st, 4th and 10th iteration, respectively. Each sample data is shown with the colour of the regressor with the highest probability. A Gaussian kernel with a kernel width of 1.0 was used to create the basis functions. Only 14 samples were retained after convergence.*

The rest of the paper is organized as follows: The algorithm for learning the one-to-many mapping using multiple RVMs is introduced in section 2. Section 3 describes a scheme for training a single RVM mapping function with multivariate outputs. The pose estimation and tracking framework is presented in section 4, and results on hand and full body tracking are shown in section 5.

1 Learning multiple RVMs

The pose of an articulated object, in our case a hand or a full human body, is represented by a parameter vector $\mathbf{x} \in \mathbb{R}^M$. The features $\mathbf{z}$ are Canny edges extracted from the image. Given a set of training examples or templates $\mathcal{V} = \{v^{(n)}\}_{n=1}^N$ consisting of pairs $v^{(n)} = \{(\mathbf{x}^{(n)}, \mathbf{z}^{(n)})\}$ of state vector and feature vector, we want to learn a one-to-many mapping from feature space to state space. We do this by learning K different regression functions, which map the input $\mathbf{z}$ to different regions in state space. We choose the following model for the regression functions

$$\mathbf{x} = \mathbf{W}^k \boldsymbol{\phi}(\mathbf{z}) + \boldsymbol{\xi}^k, \tag{4}$$

where $\boldsymbol{\xi}^k$ is a Gaussian noise vector with $\mathbf{0}$ mean and diagonal covariance matrix $\mathbf{S}^k = diag\left\{(\sigma_1^k)^2, \ldots, (\sigma_M^k)^2\right\}$. Here $\boldsymbol{\phi}(\mathbf{z})$ is a vector of basis functions of the form $\boldsymbol{\phi}(\mathbf{z}) = [1, G(\mathbf{z}, \mathbf{z}^{(1)}), G(\mathbf{z}, \mathbf{z}^{(2)}), ..., G(\mathbf{z}, \mathbf{z}^{(N)})]^T$, where G can be any function that compares two sets of image features. The weights of the basis functions are written in matrix form $\mathbf{W}^k \in \mathbb{R}^{M \times P}$ and $P = N + 1$. We use an EM type algorithm, outlined in Figure 2, to learn the parameters $\{\mathbf{W}^k, \mathbf{S}^k\}_{k=1}^K$ of the mapping functions. The regression results on a toy dataset are shown in Figure 3.

The case of ambiguous poses means that the training set contains examples

#RVMs	relevant templates	approx. total training time	mean RMS error
1	13.48 %	360 min	15.82°
5	13.04 %	150 min	7.68°
10	10.76 %	90 min	5.23°
15	9.52 %	40 min	4.69°
20	7.78 %	25 min	3.89°

(a)

(b)

Fig. 4. **(a) Single vs. multiple RVMs.** *Results of training different numbers of RVMs on the same dataset. Multiple RVMs learn sparser models, require less training time and yield a smaller estimation error.* **(b) Robustness analysis.** *Pose estimation error when using two different types of features: histograms of shape contexts (SC) and Hausdorff matching costs (HD). Plotted is the mean and standard deviation of the RMS error of three estimated pose parameters as a function of image noise level. Hausdorff features are more robust to edge noise.*

that are close or the same in feature space but are far apart in state space, see Figure 1(a). When a single RVM is trained with this data, the output states tend to average different plausible poses [1]. We therefore experimentally evaluated the effect of learning mapping functions with different numbers of RVMs (with Hausdorff fractions as the input to the mapping functions). The data was generated by random sampling from a region in the 4-dimensional state space of global rotation and scale, and projecting a 3D hand model into the image. The size of the training set was 7000 and the size of the test set was 5000. Different numbers of mapping functions were trained to obtain a one-to-many mapping from the features to the state space. The results are shown in Figure 4(a). Training multiple mapping functions reduces the estimation error and creates sparser template sets. Additionally, the total training time is reduced because the RVM training time increases quadratically with the number of data points and the samples are divided among the different RVMs.

2 Training an RVM with multivariate outputs

The RVM is a Bayesian regression framework, in which the weights of each input example are governed by a set of hyperparameters. These hyperparameters describe the posterior distribution of the weights and are estimated iteratively during training. Most hyperparameters approach infinity, causing the posterior distributions of the effectively setting the corresponding weights to zero. The remaining examples with non-zero weights are called *relevance vectors*. The attraction of the RVM is that it has good generalization perfor-

mance, while achieving sparsity in the representation. The formulation in [21] only allows regression from multivariate input to a univariate output variable. One solution is to use a single RVM for each output dimension. For example, Williams *et al.* used separate RVMs to track the four parameters of a similarity transform of an image region [24]. This solution has the drawback that one needs to keep separate sets of selected examples for each RVM. We introduce the multivariate RVM (MVRVM) which extends the RVM framework to multivariate outputs, making it a general regression tool.[1]

The data likelihood is obtained as a function of weight variables and hyperparameters. The weight variables are then analytically integrated out to a obtain marginal likelihood as function of the hyperparameters. An optimal set of hyperparameters is obtained by maximizing the marginal likelihood over the hyperparameters using a version of the fast marginal likelihood maximization algorithm [22]. The optimal weight matrix is obtained using the optimal set of hyperparameters.

The rest of this section details our proposed extension of the RVM framework to handle multivariate outputs and how this is used to minimize the cost function described in equation (1) and learn the parameters of a mapping function, $\mathbf{W}^k$ and $\mathbf{S}^k$. We can rewrite equation (1) in the following form

$$L^k = \sum_{n=1}^{N} \log \mathcal{N}(\hat{\mathbf{x}}_k^{(n)} | \mathbf{W}^k \hat{\boldsymbol{\phi}}_k(\mathbf{z}^{(n)}), \mathbf{S}^k), \tag{5}$$

$$\text{where,} \quad \hat{\mathbf{x}}_k^{(n)} = \sqrt{c_k^{(n)}} \mathbf{x}^{(n)} \quad \text{and} \quad \hat{\boldsymbol{\phi}}_k(\mathbf{z}^{(n)}) = \sqrt{c_k^{(n)}} \boldsymbol{\phi}(\mathbf{z}^{(n)}) \tag{6}$$

We need to specify a prior on the weight matrix to avoid overfitting. We follow Tipping's relevance vector approach [21] and assume a Gaussian prior for the weights of each basis function. Let $\mathbf{A} = diag(\alpha_1^{-2}, \ldots, \alpha_P^{-2})$, where each element α_j is a hyperparameter that determines the *relevance* of the associated basis function. The prior distribution over the weights is then

$$p(\mathbf{W}^k | \mathbf{A}^k) = \prod_{r=1}^{M} \prod_{j=1}^{P} \mathcal{N}(w_{rj}^k | 0, \alpha_j^{-2}), \tag{7}$$

where w_{rj}^k is the element at (r, j) of the weight matrix $\mathbf{W}^k$. We can now completely specify the parameters of the k^{th} mapping function as $\{\mathbf{W}^k, \mathbf{S}^k, \mathbf{A}^k\}$. As the form and the learning routines of parameters of each expert are the same, we drop the index k for clarity in the rest of the section. A likelihood

---

[1] Code is available from *http://mi.eng.cam.ac.uk/˜at315/MVRVM.htm*

distribution of the weight matrix $\mathbf{W}$ can be written as

$$p(\{\hat{\mathbf{x}}^{(n)}\}_{n=1}^N|\mathbf{W},\mathbf{S}) = \prod_{n=1}^N \mathcal{N}(\hat{\mathbf{x}}^{(n)}|\mathbf{W}\hat{\boldsymbol{\phi}}(\mathbf{z}^{(n)}),\mathbf{S}) \ . \tag{8}$$

Let $\mathbf{w}_r$ be the weight vector for the r^{th} component of the output vector $\mathbf{x}$, such that $\mathbf{W} = [\mathbf{w}_1,\ldots,\mathbf{w}_r,\ldots,\mathbf{w}_M]^T$ and let τ_r be the vector with the r^{th} component of all the example output vectors. Exploiting the diagonal form of $\mathbf{S}$, the likelihood can be written as a product of separate Gaussians of the weight vectors of each output dimension:

$$p(\{\hat{\mathbf{x}}^{(n)}\}_{n=1}^N|\mathbf{W},\mathbf{S}) = \prod_{r=1}^M \mathcal{N}(\tau_r|\mathbf{w}_r\hat{\boldsymbol{\Phi}}, \ \sigma_r^2) \ , \tag{9}$$

where $\hat{\boldsymbol{\Phi}} = [\mathbf{1},\hat{\boldsymbol{\phi}}(\mathbf{z}_1),\hat{\boldsymbol{\phi}}(\mathbf{z}_2),\ldots,\hat{\boldsymbol{\phi}}(\mathbf{z}_N)]$ is the *design matrix*. The prior distribution over the weights is rewritten in the following form

$$p(\mathbf{W}|\mathbf{A}) = \prod_{r=1}^M \prod_{j=1}^P \mathcal{N}(w_{rj}|0,\alpha_j^{-2}) = \prod_{r=1}^M \mathcal{N}(\mathbf{w}_r|\mathbf{0},\mathbf{A}). \tag{10}$$

Now the posterior on $\mathbf{W}$ can be written as the product of separate Gaussians for the weight vectors of each output dimension:

$$p(\mathbf{W}|\{\hat{\mathbf{x}}\}_{n=1}^N,\mathbf{S},\mathbf{A}) \propto p(\{\hat{\mathbf{x}}\}_{n=1}^N|\mathbf{W},\mathbf{S})\,p(\mathbf{W}|\mathbf{A}) \tag{11}$$

$$\propto \prod_{r=1}^M \mathcal{N}(\mathbf{w}_r|\mu_r,\boldsymbol{\Sigma}_r) \ , \tag{12}$$

where $\mu_r = \sigma_r^{-2}\boldsymbol{\Sigma}_r\boldsymbol{\Phi}^T\tau_r$ and $\boldsymbol{\Sigma}_r = (\sigma_r^{-2}\boldsymbol{\Phi}^T\boldsymbol{\Phi} + \mathbf{A})^{-1}$ are the mean and the covariance of the distribution of $\mathbf{w}_r$. Given the posterior for the weights, we can choose an optimal weight matrix if we obtain a set of hyperparameters that maximise the data likelihood in equation (12). The Gaussian form of the distribution allows us to the remove the weight variables by analytically integrating them out. Exploiting the diagonal form of $\mathbf{S}$ and $\mathbf{A}$ once more, we marginalize the data likelihood over the weights:

$$p(\{\hat{\mathbf{x}}\}_{n=1}^N|\mathbf{A},\mathbf{S}) = \int p(\{\hat{\mathbf{x}}\}_{n=1}^N|\mathbf{W},\mathbf{S})\,p(\mathbf{W}|\mathbf{A})\,d\mathbf{W} \tag{13}$$

$$= \prod_{r=1}^M \int \mathcal{N}(\tau_r|\mathbf{w}_r\hat{\boldsymbol{\Phi}},\sigma_r^2)\,\mathcal{N}(\mathbf{w}_r|\mathbf{0},\mathbf{A}) \tag{14}$$

$$= \prod_{r=1}^M |\mathbf{H}_r|^{-\frac{1}{2}}\exp(-\frac{1}{2}\tau_r^T\mathbf{H}_r^{-1}\tau_r) \ , \tag{15}$$

where $\mathbf{H}_r = \sigma_r^2\mathbf{I} + \hat{\boldsymbol{\Phi}}\mathbf{A}^{-1}\hat{\boldsymbol{\Phi}}^T$. An optimal set of hyperparameters $\{\alpha_j^{opt}\}_{j=1}^P$ and

noise parameters $\{\sigma_r^{opt}\}_{r=1}^M$ is obtained by maximising the marginal likelihood using bottom-up basis function selection as described by Tipping et al. in [22]. Again, the method was extended to handle the multivariate outputs. Details of this extension can be found in [20]. The optimal hyperparameters are then used to obtain the optimal weight matrix:

$$\mathbf{A}^{opt} = diag(\alpha_1^{opt}, \ldots, \alpha_P^{opt}) \qquad \boldsymbol{\Sigma}_r^{opt} = ((\sigma_r^{opt})^{-2} \, \hat{\boldsymbol{\Phi}}^T \hat{\boldsymbol{\Phi}} + \mathbf{A}^{opt})^{-1} \qquad (16)$$

$$\mu_r^{opt} = (\sigma_r^{opt})^{-2} \, \boldsymbol{\Sigma}_r^{opt} \, \boldsymbol{\Phi}^T \, \tau_r \qquad \mathbf{W}^{opt} = [\mu_1^{opt}, \ldots, \mu_M^{opt}]^T \qquad (17)$$

We performed experiments comparing the feature robustness in the presence of noise. This was done for features G^{HD} based on the Hausdorff distance and features based on 100-dimensional shape-context histograms G^{SC} [4]. A set of images is created by sampling a region in state space, in this case three rotation angles over a limited range, and using the sampled pose vectors to project a 3D hand model into the image. Because the Hausdorff features are neither translation nor scale invariant, additional training images of scaled and locally shifted examples are generated. After RVM training, a set of around 30 templates out of 200 are chosen for both, shape context and Hausdorff features. However note that the templates chosen by the RVM for each methods may differ. For testing, 200 poses are generated by randomly sampling the same region in parameter space and introducing different amounts of noise by introducing edges of varying length and curvature. Figure 4(b) shows the dependency of the RMS estimation error (mean and standard deviation) on the noise level. Hausdorff features are more robust to edge noise than shape context features.

3 Pose estimation and tracking

Given a candidate object location in the image we obtain K possible poses from the mapping functions, see Figure 6(a). For each mapping function $\mathbf{W}_k$ the templates selected by the RVM are matched to the input and the resulting Hausdorff fractions [11] form the basis function vector $\boldsymbol{\phi}^{HD}$. We then use regression to obtain K pose estimates via $\mathbf{x}_k = \mathbf{W}^k \boldsymbol{\phi}^{HD}$. A set of candidate object locations is obtained by skin colour detection for hands and background estimation for full human body motion. Given M candidate positions we thus obtain $K \times M$ pose hypotheses, which are used to project the 3D object model into the image and obtain image likelihoods.

The observation model for the likelihood computation is based on edge and silhouette cues. As a likelihood model for hand tracking we use the function

8

proposed in [19], which combines chamfer matching with foreground silhouette matching, where the foreground is found by skin colour segmentation. The same likelihood function is used in the full body tracking experiments, with the difference that in this case the foreground silhouette is estimated by background subtraction.

The question arises as to whether it is necessary to use the more computationally expensive model projection process for the likelihood evaluation. Being based on Gaussian regression models, RVMs do provide likelihood estimates. However, we observed that in some cases the RVM variances are too low. In order to demonstrate this we generated silhouette data from two different motions in the CMU mocap database [7], one running (173 frames) and one exercise motion sequence (4591 frames). The 62 dimensional state vector was first reduced to 6 dimensions via PCA. The RVM selected 16 training samples. When estimating the motion of the unseen exercise sequence, one expects a high uncertainty, thus large σ values for the RVM predictions, as shown in Figure 3. However, in some cases the σ values are low even though the pose is far from any pose in the training set. In other words, the RVM can be overconfident, particularly in areas of the state space that had little or no representation in the training data. Therefore, in contrast to [4] our method uses the RVM predictions only as hypotheses for a likelihood evaluation based on projecting a 3D body model into the image. It can therefore be seen as a hybrid of learning based and model based approaches.

Temporal information is needed to resolve the ambiguous poses and to obtain a smooth trajectory through the state-space after the pose estimation is done at every frame. We embed pose estimation with multiple RVMs within a probabilistic tracking framework, which involves representing and maintaining distributions of the state $\mathbf{x}$ over time.

The distributions are represented using a piecewise Gaussian model [6] with L components. The evaluation of the distribution at one time instant t involves the following steps (see Figure 6(b)):

(1) Predict each of the L components,

(2) perform RVM regression to obtain K hypotheses,

(3) evaluate likelihood computation for each hypothesis,

(4) compute the posterior distribution for each of $L \times K$ components,

(5) select L components to propagate to next time step.

The dynamics are modelled using a constant velocity model with large process noise [6], where the noise variance is set to the variance of the mapping error
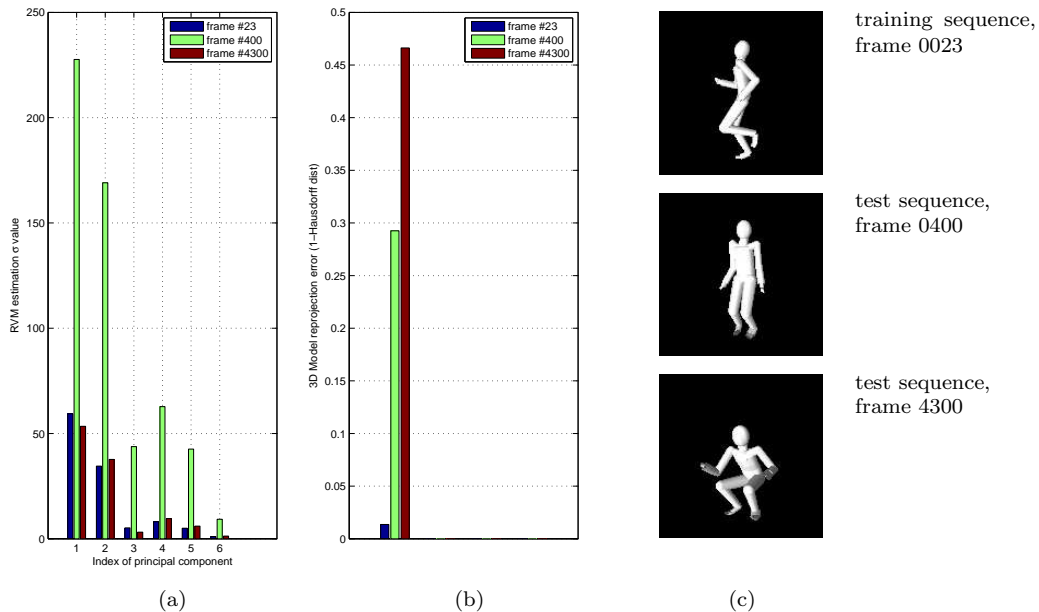
Fig. 5. **Overconfidence of RVM estimates.** *This figure shows the variance values* **(a)** *for RVM predictions for three inputs* **(c)***: one of the training sequence and two of a test sequence with previously unseen motion. We expect low uncertainty, i.e. small σ values for the training input, but large uncertainty for the unseen inputs (as for frame 400). However, in some cases the RVM is overconfident, see values for frame 4300 which are nearly identical to the values for the training input. On the other hand, the likelihood computed from the reprojection error* **(b)** *results in values that are very different from the training input.*

estimated at the RVM learning stage. At step (5) k-means clustering is used to identify the main components of the posterior distribution in the state space, similar to [23]. Components with the largest posterior probability are chosen from each cluster in turn, ensuring that not all components represent only one region of the state-space.

For a given frame the correct pose does not always have the largest posterior probability. Additionally, the uncertainty of pose estimation is larger in some regions in state space than in others, and a certain number of frames may be needed before the pose ambiguity can be resolved. The largest peak of the posterior fluctuates among different trajectories as the distribution is propagated. Hence a history of the peaks of the posterior probability needs to be considered before a consistent trajectory is found that links the peaks over time. In our experiments a batch Viterbi algorithm is used to find such a path [14].
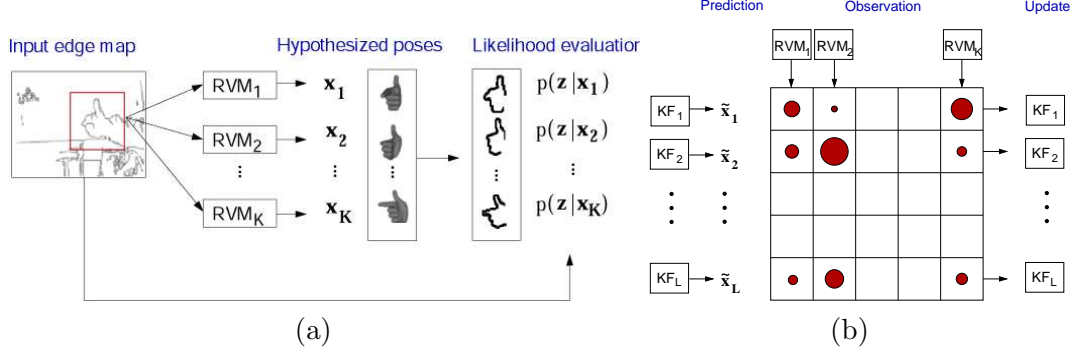
10

Fig. 6. **(a) Pose estimation.** *At each candidate location the features are obtained by Hausdorff matching and the RVMs yield pose estimates. These are used to project the 3D model and evaluate likelihoods.* **(b) Probabilistic tracking.** *The modes of likelihood distribution, obtained through the RVM mapping functions, are propagated through a bank of Kalman filters [6]. The posterior distributions are represented with an L-mode piecewise Gaussian model. At each frame, the L Kalman filter predictions and K RVM observations are combined to generate possible $L \times K$ Gaussian distributions. Out of these, L Gaussians are chosen to represent the posterior probability and propagated to the next level. The circles in the figure represent the covariance of Gaussians.*

4   Experimental results

The RVM tracking framework is applied to the problem of tracking 3D global and articulated motion of the hand and the whole human body. Throughout the experiments a Gaussian kernel with a standard deviation of 0.5 is used, a value that was found in initial experiments. The number of principal components was determined for each data set separately and was set to the minimum number of components that result in a reprojection error of less than 10%.

Full body articulation: In order to track full body motion, we use a data set from the CMU motion capture database of walking persons ($\sim$ 9000 data points). In order to reduce the RVM training time, the data is projected onto the first six principal components.

The first input sequence is a person walking fronto parallel to the camera. The global motion is mainly limited to translation. The eight-dimensional state-space is defined by two global and six articulation parameters. A set of 13,000 training samples were created by sampling the region. We use 4 RVM mapping functions to approximate the one-to-many mapping. A set of 118 relevant templates is retained after training. Background subtraction is used to remove some of the background edges. The tracking results are shown in Figure (7). The second input sequence is a video of a person walking in a circle from [17]. The range of global motion is set to 360° around axis normal to the ground plane and 20° in the tilt angle. The range of scales is 0.3 to 0.7. The

11

nine-dimensional state-space region is defined by these three global and six articulation parameters. A set of 50 000 templates is generated by sampling this region. We use 50 RVM mapping functions to approximate the one-to-many mapping. A set of 984 relevant templates is retained after training. Background subtraction is used to remove some of the background edges. The tracking results are shown in Figure (8).

Hand articulation: In this experiment we estimate the rigid body parameters as well as a lower-dimensional representation of the articulation parameters of an opening and closing hand. The method is applied to the hand sequence containing 88 frames from [19], where approximately 30 000 templates were required for tracking. To capture typical hand motion data, we use a large set of 10 dimensional joint angle data obtained from a data glove. The pose data was approximated by the first four principal components. We then projected original hand glove data into those 4 dimensions. The global motion of the hand in that sequence was limited to a certain region of the global space (80°, 60°and 40° in rotation angles and 0.6 to 0.8 in scale). The eight-dimensional state space is defined by the four global and four articulation parameters. A set of 10 000 templates is generated by random sampling in this state space. After training 10 RVMs, 455 templates out of 10 000 are retained. Due to the large amount of background clutter in the sequence, skin colour detection is used in this sequence to remove some of the background edges for this sequence. Tracking results are shown in Figure (9).

Computation time: The execution time in the experiments varies from 5 to 20 seconds per frame (on a Pentium IV, 2.1 GHz PC), depending on the number of candidate locations in each frame. The computational bottleneck is the model projection in order to compute the likelihoods (approximately 100 per second). For example, for 30 search locations and 50 RVM mapping functions result in 1500 model projections, requiring 15 seconds. It can be observed that most mapping functions do not yield high likelihoods, thus identifying them early will help to reduce the computation time.

5   Summary

This paper has introduced a framework for single camera pose estimation and tracking that is a hybrid of learning based and generative model based approaches. To this end we have developed a multivariate generalization of Tipping and Faul's [22] bottom-up method for learning a sparse RVM regressor. The method has been used as a component of a system for tracking and estimating 3D human pose from monocular image sequences. We have combined several techniques to solve this problem: (1) multivalued regression based on Gaussian mixtures to allow for multiple solutions, (2) multivariate

Fig. 7. **Tracking a person walking fronto parallel to the camera .** *The first and second rows shows the frames from [17], overlaid with the body pose corresponding to the optimal path through the posterior distribution and the corresponding the 3D model, respectively. Similarly, second and third rows show the second best path. Notice that the second path describes the walk equally well except for the right-left leg flip which is one of the common ambiguity that arises in human pose estimation from monocular view. A total of 118 templates with 4 RVM mapping functions were used.*


Fig. 8. **Tracking a person walking in a circle.** *This figure shows the results of the tracking algorithm on a sequence from [17]. Overlaid is the body pose corresponding to the optimal path through the posterior distribution, the 3D model is shown below. A total of 1429 templates with 50 RVM mapping functions were used.*

RVM for the individual regressors, (3) reprojection of a 3D model to evaluate a posteriori probabilities for the resulting 3D hypothesis, and (4) global optimization using dynamic programming to find 3D trajectories through the resulting sets of static 3D pose hypotheses.

References

[1] A. Agarwal and B. Triggs. 3D human pose from silhouettes by relevance vector regression. In *Proc. Conf. Computer Vision and Pattern Recognition*, volume II, pages 882–888, Washington, DC, July 2004.

[2] A. Agarwal and B. Triggs. Learning to track 3D human motion from silhouettes.

Fig. 9. **Tracking an opening and closing hand.** *This sequence shows tracking of opening and closing hand motion together with global motion on a sequence from [19]. A total of 537 relevant templates were used with 20 RVM mapping functions for pose estimation. As a comparison [19] used about 30 000 templates to track the same sequence.*

In *In Proceedings of the 21st International Conference on Machine Learning*, pages 9–16, Banff, Canada, 2004.

[3] A. Agarwal and B. Triggs. Monocular human motion capture with a mixture of regressors. In *In IEEE Workshop on Vision for Human Computer Interaction*, 2005.

[4] A. Agarwal and B. Triggs. Recovering 3D Human Pose from Monocular Images. In *Trans. PAMI*, vol. 28, no. 1, pages 44-58, January 2006.

[5] M. Brand. Shadow puppetry. In *Proc. 7th Int. Conf. on Computer Vision*, volume II, pages 1237–1244, Corfu, Greece, September 1999.

[6] T. J. Cham and J. M. Rehg. A multiple hypothesis approach to figure tracking. In *Proc. Conf. Computer Vision and Pattern Recognition*, volume II, pages 239–245, Fort Collins, CO, June 1999.

[7] Carnegie-Mellon MoCap Database. http://mocap.cs.cmu.edu

[8] P. Felzenszwalb, D. Huttenlocher Pictorial Structures for Object Recognition. In *International Journal of Computer Vision*, Vol. 61, No. 1, January 2005.

[9] D. M. Gavrila. Pedestrian detection from a moving vehicle. In *Proc. 6th European Conf. on Computer Vision*, volume II, pages 37–49, Dublin, Ireland, June/July 2000.

[10] N. R. Howe, M. E. Leventon, and W. T. Freeman. Bayesian reconstruction of 3D human motion from single-camera video. In *Adv. Neural Information Processing Systems*, pages 820–826, Denver, CO, November 1999.

[11] D. P. Huttenlocher, J. J. Noh, and W. J. Rucklidge. Tracking non-rigid objects in complex scenes. In *Proc. 4th Int. Conf. on Computer Vision*, pages 93–101, Berlin, May 1993.

[12] M. W. Lee, I. Cohen. Human upper body pose estimation in static images. In *Proc. 8th European Conf. on Computer Vision*, pages II: 126–138, 2004.

[13] M. Jordan, R. Jacobs. Hierarchical mixtures of experts and the EM algorithm. *Neural Computation*, 6:181–214, 1994.

14

[14] R. Navaratnam, A. Thayananthan, P. H. S. Torr, R. Cipolla. Hierarchical Part-Based Human Body Pose Estimation In *Proc. British Machine Vision Conference*, London, UK, 2005.

[15] R. Rosales, V. Athitsos, L. Sigal, and S. Scarloff. 3D hand pose reconstruction using specialized mappings. In *Proc. 8th Int. Conf. on Computer Vision*, volume I, pages 378–385, Vancouver, Canada, July 2001.

[16] G. Shakhnarovich, P. Viola, and T. Darrell. Fast pose estimation with parameter-sensitive hashing. In *Proc. 9th Int. Conf. on Computer Vision*, volume II, pages 750–757, 2003.

[17] H. Sidenbladh, F. De la Torre, and M. J. Black. A framework for modeling the appearance of 3D articulated figures. In *IEEE International Conf. on Automatic Face and Gesture Recognition*, pages 368–375, Grenoble, France, 2000.

[18] C. Sminchisescu, A. Kanaujia, Z. Li, and D. Metaxas. Discriminative density propagation for 3D human motion estimation. In *Proc. Conf. Computer Vision and Pattern Recognition*, pages 217–323, June 2005.

[19] B. Stenger, A. Thayananthan, P. H. S. Torr, and R. Cipolla. Model-Based Hand Tracking Using a Hierarchical Bayesian Filter. In *Trans. PAMI*, vol. 28, no. 9, pages 1372-1384, September, 2006.

[20] A. Thayananthan. *Template-based pose estimation and tracking of 3D hand motion*. PhD thesis, University of Cambridge, UK, 2005.

[21] M. E. Tipping. Sparse Bayesian learning and the relevance vector machine. *J. Machine Learning Research*, pages 211–244, 2001.

[22] M. E. Tipping and A. Faul. Fast marginal likelihood maximisation for sparse bayesian models. In *Proc. Ninth Intl. Workshop on Artificial Intelligence and Statistics*, Key West, FL, January 2003.

[23] J. Vermaak, A. Doucet, and P. Pérez. Maintaining multi-modality through mixture tracking. In *Proc. 9th Int. Conf. on Computer Vision*, 2003.

[24] O. Williams, A. Blake, and R. Cipolla. A sparse probabilistic learning algorithm for real-time tracking. In *Proc. 9th Int. Conf. on Computer Vision*, volume I, pages 353–360, Nice, France, October 2003.